AFRL-IF-RS-TR-2006-122
Final Technical Report
April 2006

RANDOM ARRAY THROUGH THE WALL IMAGING SENSOR

AKELA, Incorporated

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2006-122 has been reviewed and is approved for publication

APPROVED: /s/

PETER J. COSTIANES
Project Engineer

FOR THE DIRECTOR: /s/

JOSEPH CAMERA, Chief
Information & Intelligence Exploitation Division
Information Directorate

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	APRIL 2006	Final Apr 04 – Oct 05

4. TITLE AND SUBTITLE
RANDOM ARRAY THROUGH THE WALL IMAGING SENSOR

5. FUNDING NUMBERS
C - F30602-03-C-0085
PE - N/A
PR - TWIS
TA - 03
WU - 01

6. AUTHOR(S)
Alan Hunt

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
AKELA, Incorporated
5551 Ekwill Street, Suite A
Santa Barbara California 93111

8. PERFORMING ORGANIZATION REPORT NUMBER
N/A

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Research Laboratory/IFEC
525 Brooks Road
Rome New York 13441-4514

10. SPONSORING / MONITORING AGENCY REPORT NUMBER
AFRL-IF-RS-TR-2006-122

11. SUPPLEMENTARY NOTES

AFRL Project Engineer: Peter J. Costianes/IFEC/(315) 330-4030/ Peter.Costianes@rl.af.mil

12a. DISTRIBUTION / AVAILABILITY STATEMENT
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

12b. DISTRIBUTION CODE

13. ABSTRACT ***(Maximum 200 Words)***

The activities performed by AKELA under contract F30602-03-C-0085 have been designed to meet the operational need for better tactical surveillance information for military and law enforcement personnel while conducting urban operations. We have developed a small, robust, cost effective radar capable of detecting both large and small motions at substantial standoff distances through steel reinforced concrete walls. The software we have developed to control the radar and perform imaging has demonstrated the capability to perform imaging in real time at frame rates up to 5 frames a second. We have characterized in detail the frequency dependent attenuation of most common wall types. The system has demonstrated the ability to be used with antennas in random locations using wireless control. We have made detailed measurements and predictions of the system's probability of detecting individuals through different types of walls as a function of distance. And we have demonstrated the ability for sensor self location. The data collected and analysis performed this program are sufficient to allow point designs for STTW systems to be made.

14. SUBJECT TERMS
Stepped Frequency Radar, Imaging Radar, Through Wall Imaging

15. NUMBER OF PAGES
59

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL

Table of Contents

List of Figures

List of Tables

Introduction

While early development of through-the-wall detection and imaging systems was supported by the intelligence community, it is only through the changing nature of recent regional conflicts and the advent of a more precision method of warfare that the military has begun to focus on similar systems for urban operations. Technology development in this area has primarily been oriented toward law enforcement and peacekeeping operations and supported mainly by the National Institute of Justice.

There are many situations in both peace keeping and law enforcement operations where there is a need to not only determine if there is someone inside a building structure but also to know where they are. These situations arise during searches for suspects, hostage and barricade incidents, and tactical surveillance. While in many cases the objective is to make contact with a suspect to defuse a potentially violent situation, not all suspects are cooperative. This often leads to the necessity of entering a building in order to take a suspect into custody. Unfortunately, most operations conclude with a physical search with personnel under a great deal of stress and subject to a high possibility of physical harm. Military operations in urban terrain face this same set of issues, but greatly magnified by the scope of operations and the presence of unarmed, noncombatant, civilians.

The technology available to military and law enforcement personnel to assist them in securing a building is rapidly changing. A new generation of techniques is being developed that will enhance the capability to look into buildings at standoff distances and build a picture of the tactical situation. Knowing that there is someone inside a building, where they are inside, and what the internal layout of the building is, will change the operational tactics used and increase the probability that an operation will successfully conclude without casualties. While these new systems work, they have a number of operational deficiencies.

High frequency motion detection systems require that the detection equipment be held steady enough that pressing against a wall is the expected way to use the device. This gives an indication of human presence or other moving object, but little other information. Hand held imaging radars give a better picture of what is behind a wall but have the same requirement for holding against a wall. They do not give real time continuous data, and the data that they do give is available only to the operator. Resolution is poor for this type of system because its small size limits the imaging resolution and the number of targets that can be individually resolved.

A more robust and useful system would involve a distributed imaging radar. In this type of system each person involved in a building clearing operation would carry a small sensor. When the signals from all of the sensors were combined with one another, a continuous image of the area being cleared would be available to each of the individuals as well as their command elements. Commanders could then more efficiently control the deployment of forces and be aware of the situation inside the building being cleared. Such a

system would locate and identify both the clearing force and hostile forces. With each individual carrying only part of the sensor system, there is no one point of failure and loss of sensor nodes would result in a graceful degradation of performance. The ability to resolve two individuals close to one another would be improved because of the greater separation between sensors than in a handheld system.

The activities performed by AKELA under contract F30602-03-C-0085 have been designed to meet this operational need for better tactical surveillance information. We have developed a small, robust, cost effective radar. We have demonstrated the capability to detect both large and small motions with the radar at substantial standoff distances through steel reinforced concrete walls. The software we have developed to control the radar and perform imaging has demonstrated the capability to perform imaging in real time at frame rates up to 5 frames a second. We have characterized in detail the frequency dependent attenuation of most common wall types. The system has demonstrated the ability to be used with antennas in random locations using wireless control. We have made detailed measurements and predictions of the system's probability of detecting individuals through different types of walls as a function of distance. And we have demonstrated the ability for sensor self location. Figure 1 is a picture of the brassboard radar sensor that was developed for this program.

This radar was used to perform extensive static and dynamic testing. Static tests were performed with antennas placed varying distances from the walls to gather data that was used to characterize the electro-

Figure 1 - Independent radar sensor units.

magnetic material properties of the walls. Static testing in an open field environment provided data that was used to theoretically determine the maximum detectability range of our imaging radar system and to determine the sensitivity of different motion detection algorithms.

Analysis of the static data resulted in the generation of frequency dependent attenuation and dielectric constant curves for each of the wall types. Figure 2 shows a comparison of the frequency dependent attenuation characteristics for five different types of walls we measured over the band of .5 to 2 GHz. It was also found that there is very little frequency dependence of the dielectric constant.

Dynamic testing was also performed to gather data that was used to develop and test both large and small motion detection algorithms through walls and in an open field environment. It was found that both large and small motions could be detected at significant ranges through all types of walls as shown in Table 1. Open field testing provided us with data that could be used to predict the limits of system detectability. Figure 3 shows curves of probability of detection versus distance for various wall types. As shown in the figure, the radar we have developed has the sensitivity to detect an individual through 12" of steel reinforced concrete at a distance of 20 meters and for other wall types at distances in excess of this number. While this falls short of our desired goal of 30 meters, it is none the less a significant result.

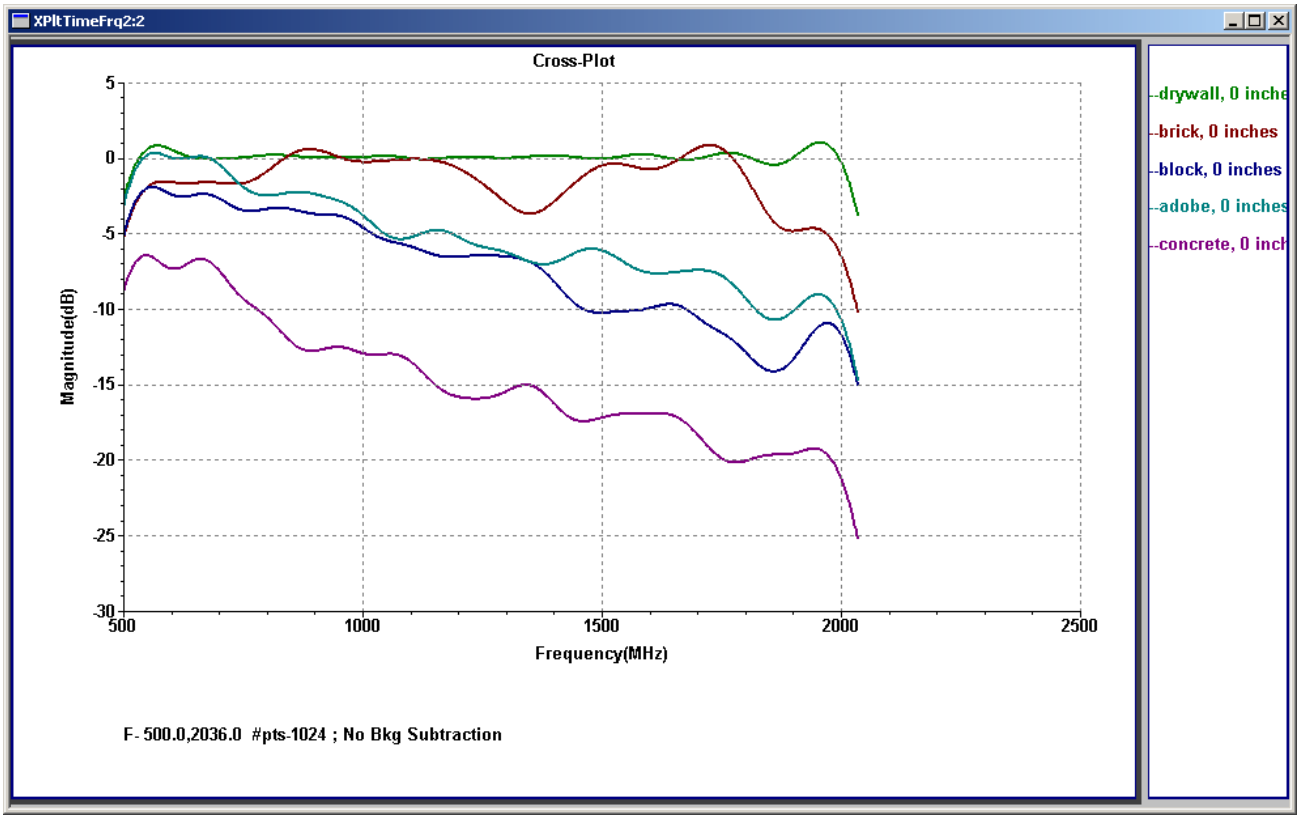

Figure 2 - Attenuation characteristics of wall materials.

Drywall: at 19´ from standoff of 12´
Brick: at 21´ from standoff of 14´
Concrete block: at 28´ from standoff of 21´
Adobe: at 16´ from standoff of 4´
Reinforced concrete: at 22´ from standoff of 10´

Table 1 - Distances of detection for breathing and large motion for various walls.

Finally, we demonstrated the capability to detect large and small motions in real time. The software that we developed allowed us to use an Ethernet interface to control the radar and collect data at its maximum rate of 4.096 Mbps. The motion detection algorithms were implemented in the program in a manner that allowed us to toggle between them at any time to allow comparison of the different algorithms under different conditions. Using this software, we were able to demonstrate real time operation of the radar using a

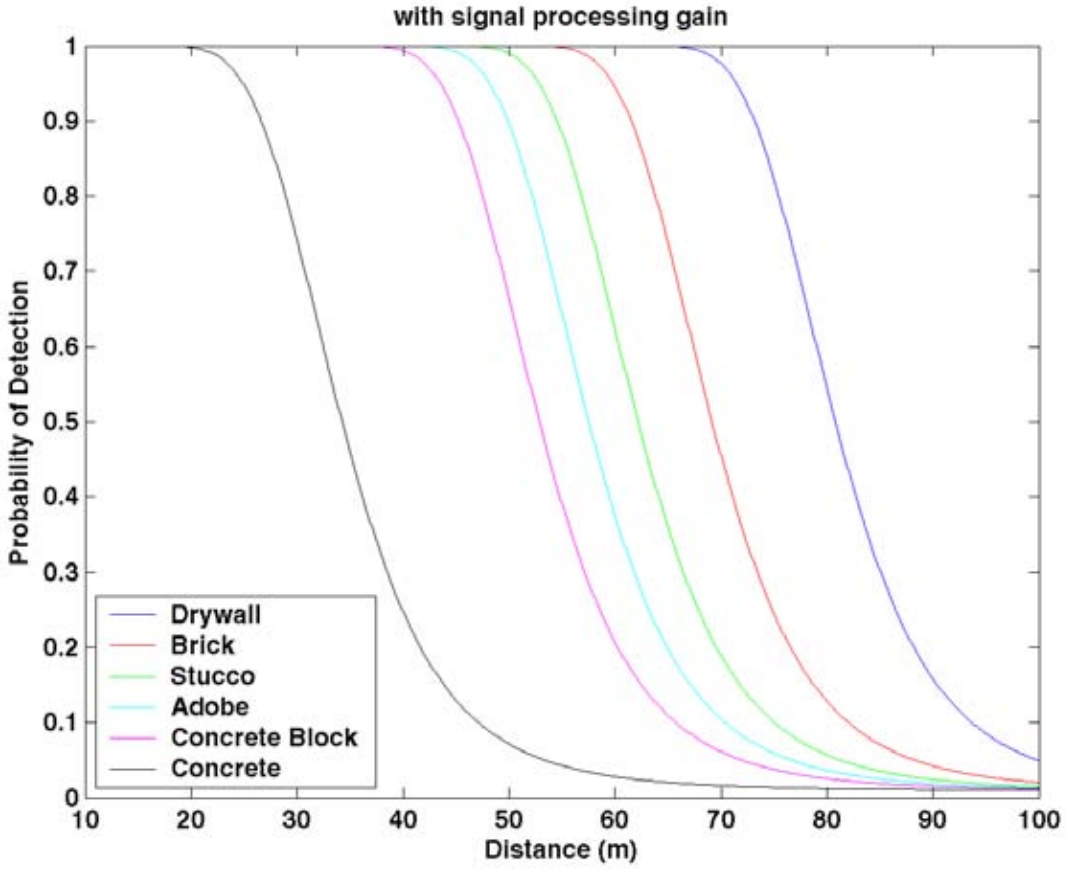

Figure 3 - Probability of detection through various wall materials using motion detection.

2.8 GHz computer with a hardware graphics accelerator card. Depending on the number of data points in a sweep and the area over which an image is reconstructed, we produce images at frame rates of between 6 and 10 frames per second. On a laptop computer with the same speed processor but no graphics accelerator, the demonstrated frame rate is between 3 and 4 frames per second. Figure 4 shows an image captured from a real time demonstration at AKELA of a person's breathing response through an interior and exterior wall combination at a distance of 6 meters.

While imaging with independent units in random locations was not as clear as anticipated due to noise sources in the radar and limitations in the image processing algorithms, the objectives of our program have been met. We have developed a sensitive, cost effective radar sensor, demonstrated real time detection of individuals through challenging wall materials, and developed an understanding of the issues that need to be resolved to improve imaging with independent radar sensors. The development performed, data collected, and analysis performed this program are sufficient to allow point designs for through the wall systems to be made.

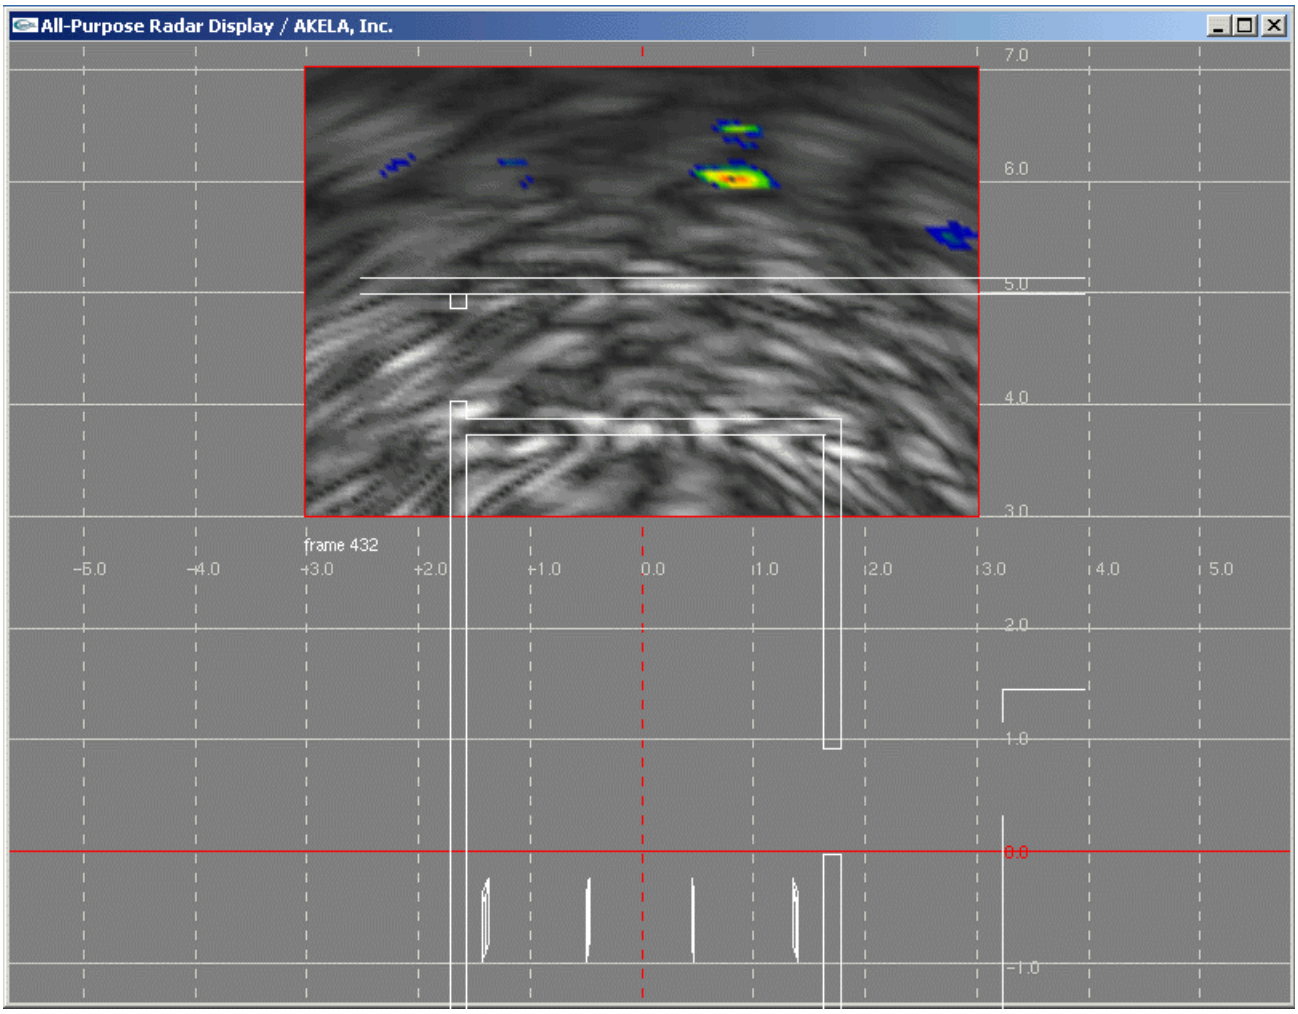

Figure 4 - Breathing response of person through two walls.

System Overview

Sensors that are designed to create images of the interior of an object most often use tomographic image reconstruction. While the type of energy used to probe the object varies, the basic method of reconstructing an image of the inside of the object consists of collecting data over a range of orientations and using backward propagation algorithms to remove ambiguities in the location of objects that scatter the probing radiation. In this way, a picture of the inside of the object showing the relative location of scattering objects is formed.

The imaging system we developed uses tomographic image reconstruction and has been designed with three major subsystems. They are 1) a digitally controlled, stepped frequency radar operating over the band from 500 MHz to 2 GHz, 2) an antenna array, and 3) a system control, processing, and display computer. The radar is the source of probing radiation, the antenna array serves as the means of varying the probing radiation over the field of view of the sensor system, and the computer provides the platform for solving the mathematical algorithms that use the scattered radiation to determine the location of objects that scatter the probing radiation. Each plays an important role in the success of the imaging system.

Radar Subsystem

In order to create an image of the interior of an object, the illuminating radiation must be able to pass through the object with little attenuation. For through the wall surveillance systems the material properties of the wall determine the degree to which the system will be successful. The major considerations are the absorption and refraction losses for the penetrating radiation. Data taken by Frazier[1] show that most building materials are relatively transparent below 4 GHz. A graph that summarizes the one way attenuation loss through different types of materials is shown in Figure 5. For the case of reinforced concrete walls, Frazier made measurements at frequencies between 100 MHz and 2.5 GHz, finding that attenuation was moderate up to 2 GHz before beginning to increase rapidly. Based on this data, a system that uses frequencies below 2 GHz has the best chance of seeing through walls.

The operating range of our radar, 500 MHz to 2 GHz, was selected on the basis of both technical and operational considerations. Low frequency cutoff was determined primarily by the size of the antenna needed to radiate efficiently at that frequency. A half wavelength at 500 MHz is 0.3 m (≈ 12”) and at 250 MHz is 0.6 m (≈ 24”). A large antenna is impractical for a surveillance system that is supposed to be portable. High frequency cutoff was determined primarily by the availability of modestly priced electronic components. Selecting higher operating frequencies makes the design much more difficult and costly. We get reasonably good penetration of common building materials then, with a system that isn’t too big to be useful, or too costly to own.

The specific radar design is CW, bistatic, and stepped frequency. CW operation means that the radar sends and receives at the same time. Bistatic means that there are separate transmit and receive antennas. Stepped frequency means that the radar only transmits a single frequency at a time, but that it has the ability to transmit over a wide range of frequencies by stepping through them one at a time.

CW operation gives the radar certain technical advantages. Since each transmitted frequency is known precisely, it is possible to use narrowband detection techniques to improve signal to noise ratio and improve the ability to reject signals in adjacent frequency bands (jam immunity). A sample of the transmitted signal is sent to the receiver and used as a phase reference for the received signal. This allows the radar to demodulate the received signal into its in phase and quadrature components and results in a coherent system. That means that from frequency to frequency and sweep to sweep, the received signals from all of the stationary objects in the scene can be added to improve system signal to noise performance.

Bistatic operation simplifies the radar design. It is difficult to use the same antenna to simultaneously transmit and receive because of leakage between the transmitter and receiver. Receiver sensitivity is set by the level of the transmitted signal that gets directly into the receiver. Its effect is not unlike that of having someone standing next to you talking loudly while you are trying to hear what someone else is saying across the room. Providing sufficient electronic isolation over a wide bandwidth is very difficult, so the easiest method (as is also the case for the conversation situation above) is to separate the transmit and receive antennas.

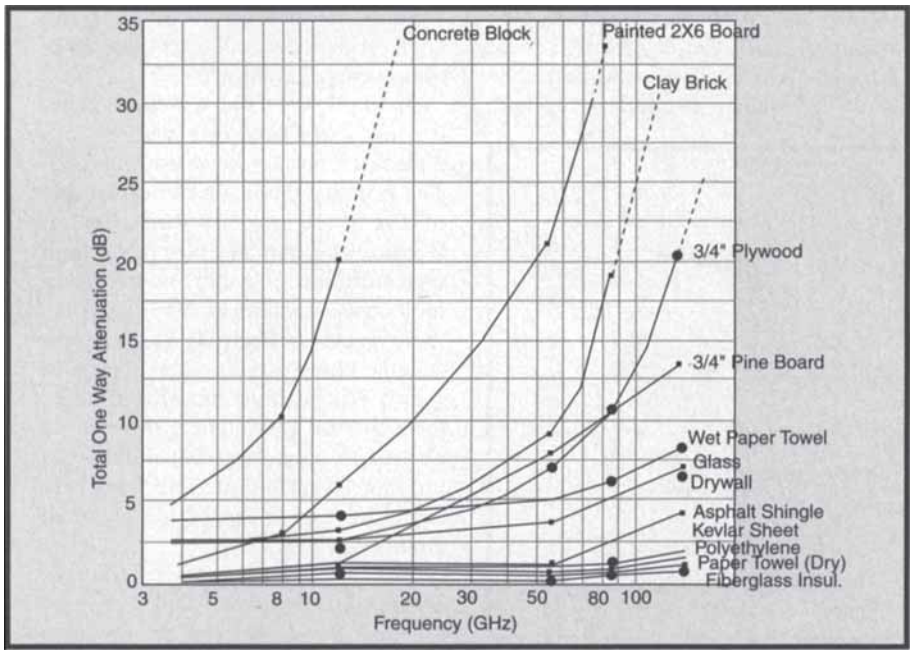

Figure 5 - Attenuation of common building materials. [1]

1. L. Frazier, "MDR for Law Enforcement", IEEE Potentials, Vol. 16, No. 5, pp. 23 - 26, 1998

Stepped frequency operation also provides technical advantages. Being able to select frequencies gives the radar the flexibility to change its range resolution, avoid transmitting on critical communication or navigation frequencies, and to optimize its waveform to enhance the performance of signal processing algorithms. For a stepped frequency waveform, range resolution is determined by the bandwidth of the waveform and follows the relation $\Delta R = c/2n\Delta f$, where c is the speed of light, n the number of frequency steps in the waveform, and Δf is the frequency step size. As with pulsed radar waveforms, there is range ambiguity associated with a stepped frequency waveform. The unambiguous range limit $R_u = c/2\Delta f$. Fourier transforming the collected spectral information produces a range profile. Unfortunately, scatterers at ranges greater than the range limit will fold over and appear in the range profile.

For our radar system, the total bandwidth is 1.5 GHz, making the size of a resolvable range cell 0.1 meters (3.9 inches). The unambiguous range limit for a waveform of 1024 points at this bandwidth is 102.4 meters. Reducing the number of points over the same bandwidth decreases the range limit. Reducing the bandwidth but keeping the same number of points increases the range limit, but also increases the size of a resolvable range cell.

Antenna Array Subsystem

To achieve two dimensional image reconstruction requires being able to resolve scattering objects in the cross range direction as well. Resolution in the cross range direction can only be achieved by varying the illumination over the field of view of the sensor system. Figure 6 shows the situation of a sensor system with a wide beamwidth. The sensor sees contributions from all cross range elements that are at the same dis-

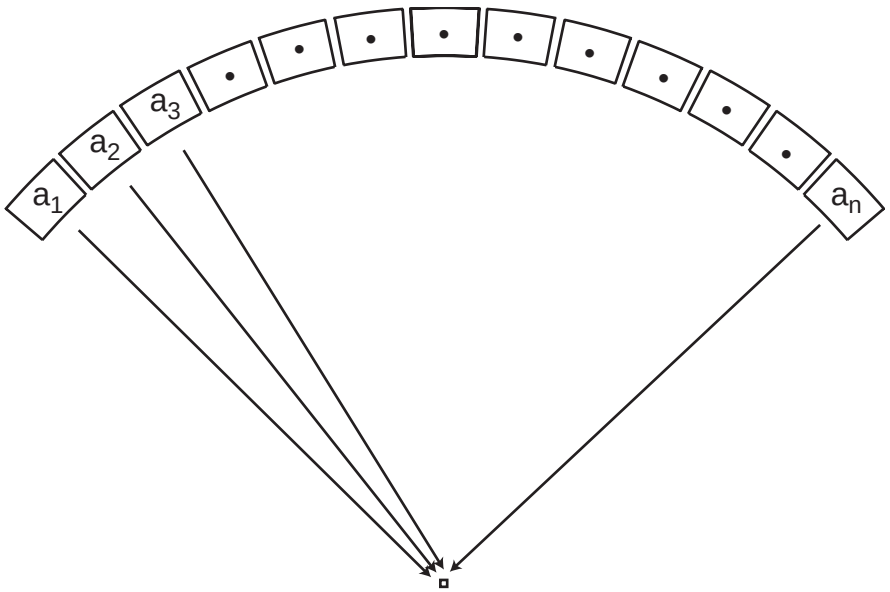

Figure 6 - Each cross range cell contributes to the signal collected at the sensor.

tance from the sensor. It is possible to resolve objects very precisely in the range direction. However, because the beamwidth of the antenna is large, the cross range location of the object can be anywhere along the constant range arc shown in the figure.

To improve resolution in the cross range direction it is necessary to produce a set of linearly independent observations in the cross range direction sufficient to mathematically solve for the signal contribution of each cross range cell.

As shown in the figure the rectangular elements, a_n, lie at a constant range from the sensor. For the case shown, these points are all on an arc whose width is the width of the sensor beam pattern. In a stepped frequency radar system, when we transform the frequency sweep to the time domain, each point in time represents a specific distance from the antenna. The signal sensed by the radar is the sum of the signal reflections from each point along the constant distance arc.

The preferred method of improving the cross range resolution is by use of a very narrow antenna beam that can be swept over the field of view. At the frequencies of operation where there is good penetration of building materials, antennas with very narrow beam widths are too large to be practical. The common practice for radar image reconstruction in cases where increasing the size of the antenna is impractical, is to move the radar taking data at various intervals, and then to synthesize an antenna aperture to obtain cross range resolution.

Rather than move the radar, it is also possible to keep the radar at a single location and synthesize an aperture from an array of antennas. An array where all antennas receive simultaneously and the illumination is provided by a single antenna is called a real array. An array that is scanned from end to end, transmitting with only a single antenna at a time, is called a synthetic array. Figure 7 shows a comparison of the antenna patterns for a single antenna, four antennas used as a real array, and four antennas used as a synthetic array. The array is made up of antennas with patterns that are the same as that of the single antenna. The decrease in beamwidth achieved by scanning the array is noticeable, and it improves the ability of the system to resolve scatterers in the cross range direction.

While resolution is constant in the range direction, it degrades with distance in the cross range direction. This is because the width of the antenna beam naturally diverges with distance. The cross range resolution is a function of the wavelength at the lowest operating frequency of the radar, the length of the physical antenna aperture, and the distance to the target being imaged. Figure 8 shows schematically the system parameters that determine cross range resolution of a radar system that scans an antenna array to synthesize an aperture. Figure 9 shows a set of charts that illustrate the system dependence of cross range resolution. The cross range resolution shown is a measure of how close two targets can be to one another and still be resolved as two targets by the radar.

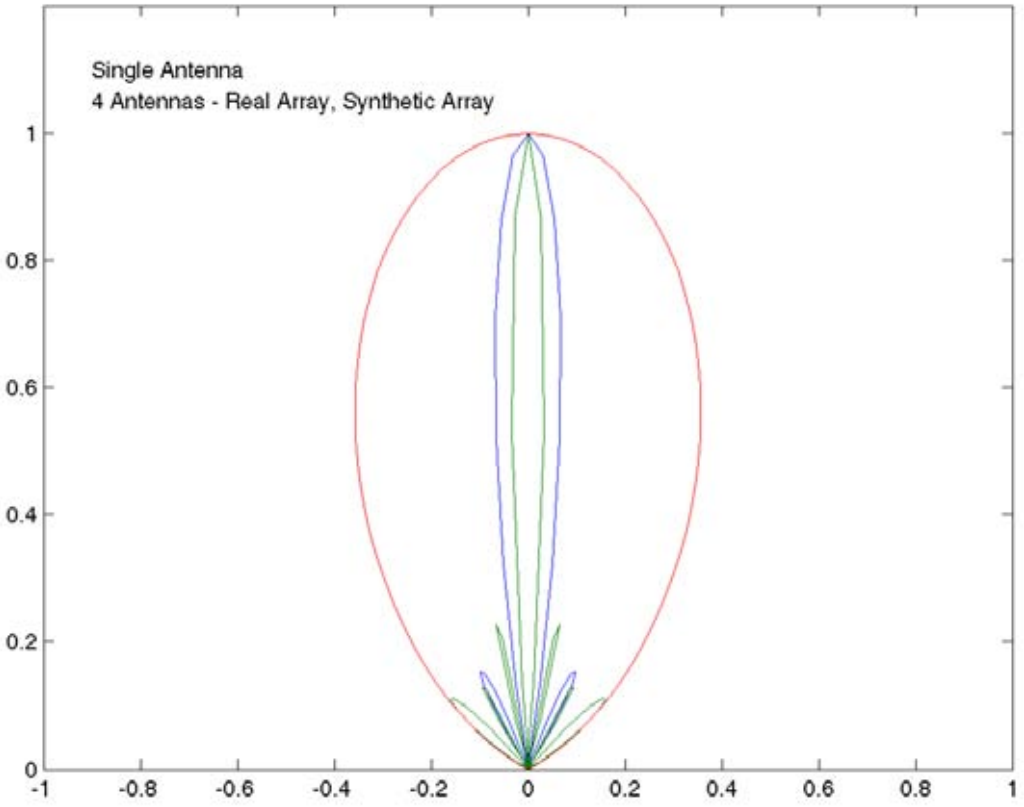

Figure 7 - Scanning a multiple element array improves cross range resolution.

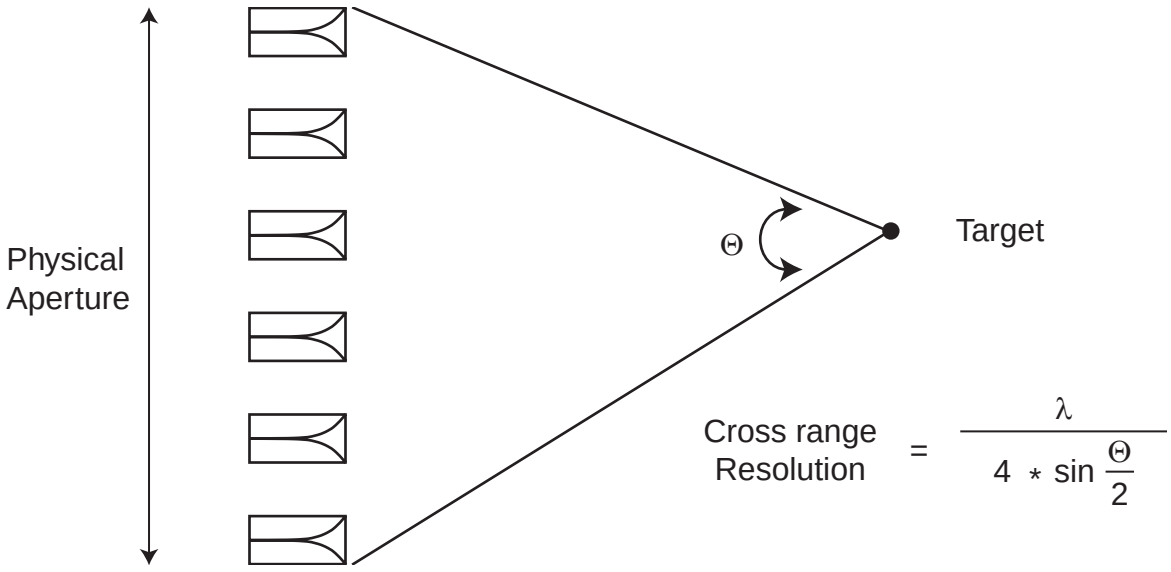

Figure 8 - Cross range resolution of a scanned array radar system.

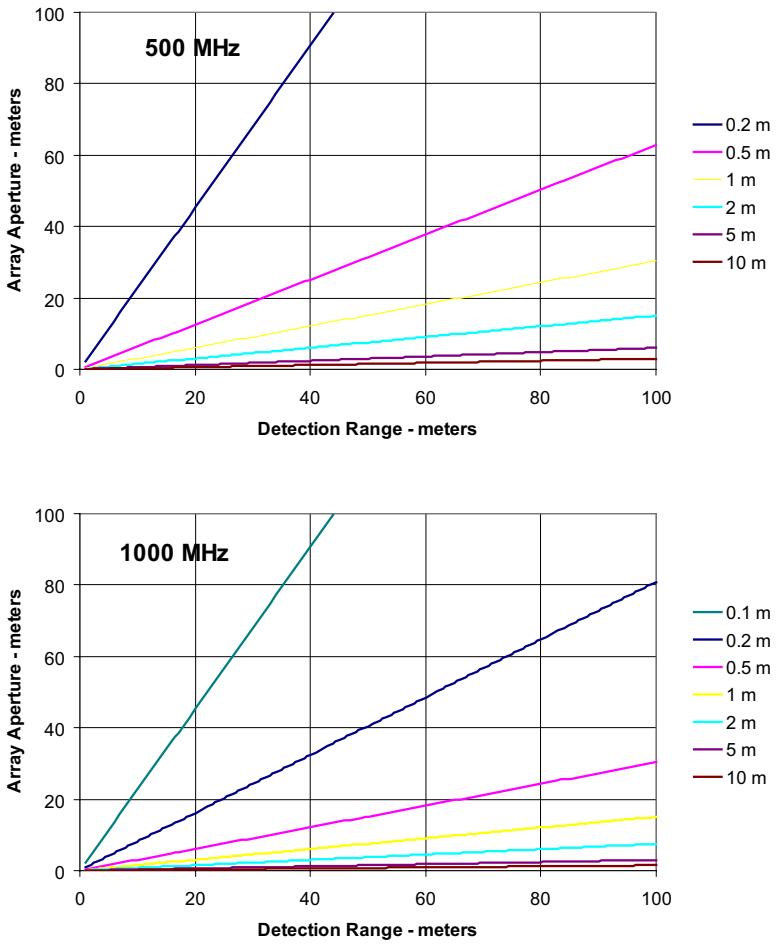

Figure 9 - Cross range resolution depends on system operating parameters.

Control, Processing, and Display Subsystem

This subsystem controls the operating parameters of the radar, conditions the received data, performs the image reconstruction, and displays the results.

Image reconstruction is accomplished by either of two methods - Fourier transform or backward propagation. Figure 10 shows schematically the Fourier transform method. Taking the Fourier transform of the stepped frequency data for an individual transmit/receive antenna pair creates a range profile for all of the scatterers in the antenna field of view. The bistatic range to each pixel in the image map is used to index into the range profile to find the value of the in phase and quadrature components of the scattered field from that range. For a different antenna pair, the bistatic range to the same pixel will be different. The values from all of the antenna pairs from the scanned array are summed for each pixel in the image map. Where there are objects in the image that result in scattering, the individual observations from the antennas will be in phase and sum to a large value. Where there are no objects, the individual observations will be out of phase and

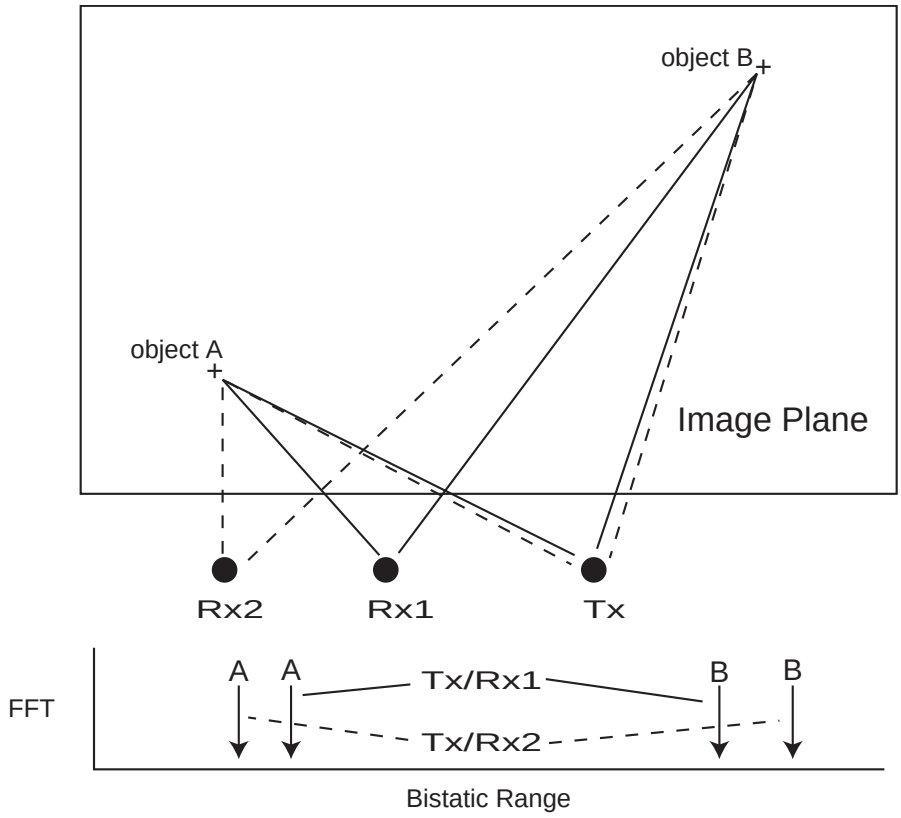

Figure 10 - Image reconstruction using the FFT.

tend to sum toward zero. The magnitude of the summation depends on the radar cross section of the scattering object and the distance from the antenna array. Figure 11 shows a schematic view of the reconstruction of a single point in the image map for a 4 element array.

The FFT method has the advantage of being relatively fast computationally. Distance weighting can be applied to the result to account for the distance dependence of the magnitude of the electric field, however, frequency dependent effects such as antenna pattern variations can not be accounted for.

In the backward propagation method shown in Figure 12, for each pixel in the image, the data is phase adjusted a single frequency at a time to account for the change in phase associated with the bistatic distance between the transmitter and receiver. During phase adjustment, the effects of range attenuation and frequency dependent antenna patterns can be easily included. As in the FFT method, the contributions at each pixel from each antenna pair are summed to create a subimage at each frequency. The final image is formed by summing all of the subimages. In addition to the ease with which corrections for antenna pattern effects can be made, the backward propagation method avoids the necessity of using interpolation as the FFT must do when the bistatic range to a pixel and the FFT range index are not exactly the same. The major disadvantage of the backward propagation method is that it is more computationally intensive.

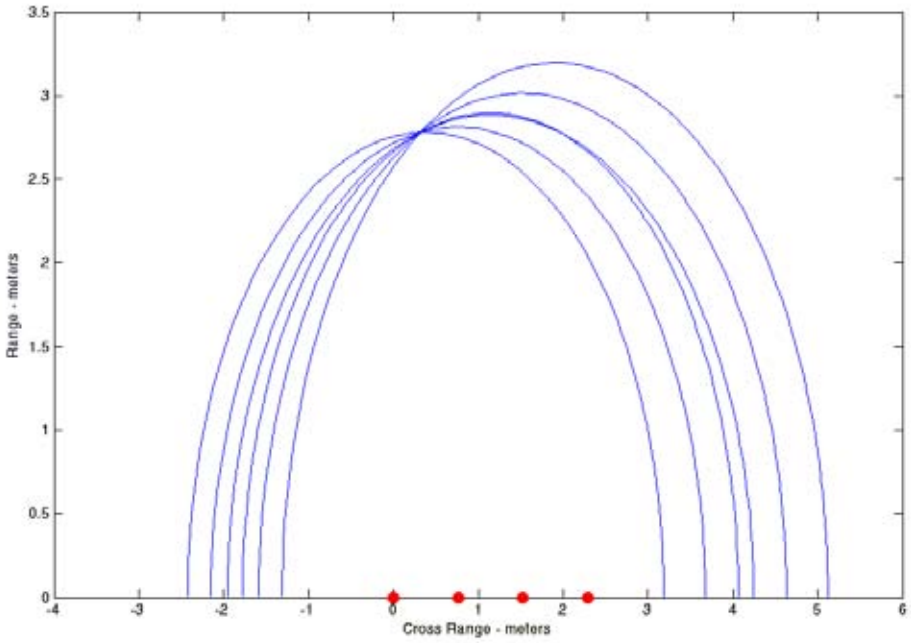

Figure 11 - Reconstruction of a point in the image map.

For cases where there are many objects that scatter energy within the field of view of the antenna array, there will be places in the image map where there will be ghost targets. These targets show up at places where the range profiles of a subset of the antenna pairs intersect. An example is the combination of a multipath return from one antenna pair, coinciding with the direct path return from another scatterer in the image plane. In these cases, the signal contributions sum, but generally to a lower value than for actual targets. The difference between real and ghost targets becomes larger as the signal to noise ratio of the image improves.

Increasing signal to noise ratio can be accomplished by averaging image frames. If Y(n) is the updated full spectrum image, and X(n) is the n^{th} full spectrum image, then Y(n) = (1/n)X(n) + (n-1/n)Y(n-1). An important assumption in applying this scheme is that there are not significant changes from image to image. This updating process is stable and the sensitivity to target motion is not high because the most current image is not weighted heavily. If there are moving targets, they will leave a trail in the image.

Another method for reducing the effect of ghosts is to use additional processing within individual images. Two common techniques are moving average filtering and median filtering. Moving average filters are low pass filters and result in trade-offs between resolution and smoothness of the image. A median filter applies a weighting function that varies inversely as the standard deviation of the individual antenna observations at pixels associated with potential targets. Observations for direct path scatterers tend to have very small standard deviations while those for ghosts have a much wider variation. The median filter is not a linear algorithm and is relatively complicated to implement for complex value images.

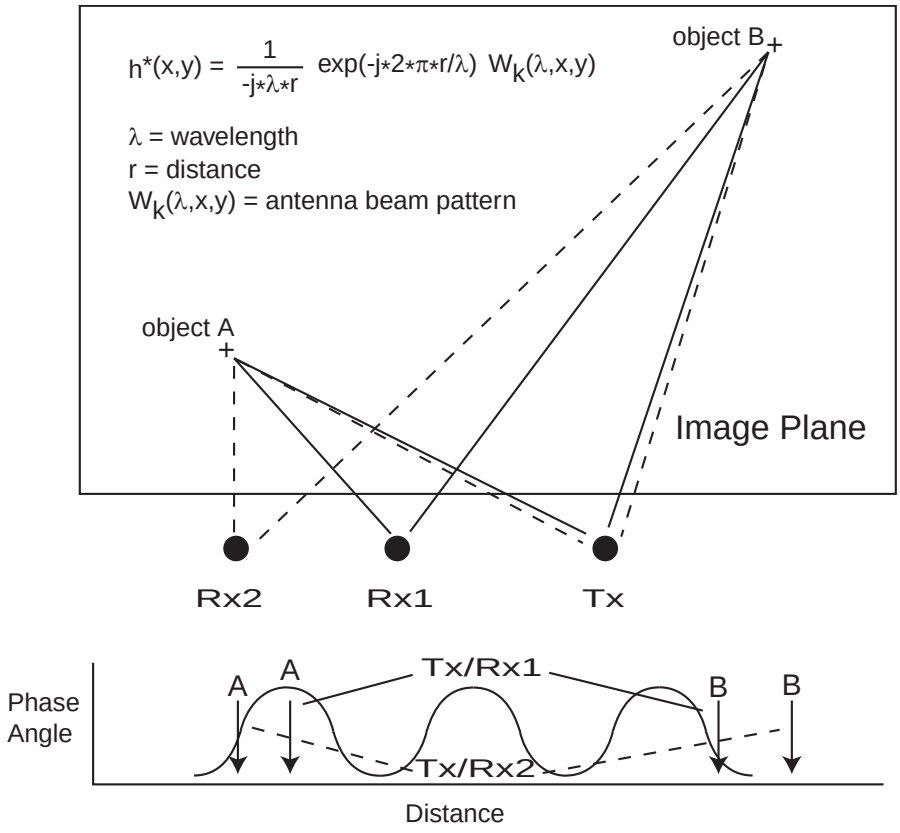

Figure 12 - Image reconstruction using backward propagation.

Motion detection is performed using complete image frames. In its simplest implementation, this is accomplished by subtraction of adjacent images in a sequence. If X(n) is the full-spectrum image sequence, then the 1st order motion detection image Z(n) = X(n) - X(n-1). Higher order detectors are more stable with better signal to noise ratio, but they are less sensitive. The 2nd order motion detection image can be expressed as Z(n) = (0.5)X(n) - X(n-1) + (0.5)X(n-2).

It is possible to superimpose motion profiles on the underlying image. The common approach is to form the display image D(n) = (a)abs[Z(n)] + (1-*a*)abs[Y(n)] where *a* is the parameter, bounded between 0 and 1, that controls the threshold for visualization. Larger *a* emphasizes the changes due to motion with relatively light background profiles. When *a* = 1, only motion is displayed. In order to reduce confusion, it is often the case that color is used to represent the motion term Z(n), and grey scale is used to represent the background image Y(n).

Program Point of Departure

Testing of our first generation brassboard imaging system developed on contract F30602-00-C-0205 showed several areas where there were performance weaknesses. The major weakness of the radar subsystem was the frequency stepping speed of the radar. This weakness shows up primarily in the ability

of the system to account for movement. Since we use a coherent addition process to form images, anything that moves during the acquisition of an imaging frame adds noise to the image. This makes signal averaging to achieve a better signal to noise ratio much less effective. It also makes it difficult to detect fast moving targets.

For the antenna array subsystem, we found that the fixed, linear configuration limited cross range resolution and exacerbated the effects of ghosts in the images. Figure 13 shows this limitation. The image to the left in the figure is of a stud wall. The wall is parallel to the boresite of the imaging radar array and lined up approximately with the left edge of the array. Each of the studs is clearly evident in the image. However, note that there are also three other similar patterns in the image. A simulation of the image reconstruction technique shows that the ghosts are associated only with the geometry of the linear array and the spacing of the studs in the wall. The results of the simulation are shown to the right in the figure. There are ghost intersections in the same locations as those that show up in the experimental data and there are no frequency dependent effects.

As indicated earlier, one method of reducing the effect of ghosts is to use additional processing within individual images. An easier method is to use a random spacing for the array of antennas. The image reconstruction algorithms do not require an even spacing. Figure 14 shows an image reconstruction simulation that demonstrates the effect of random spacing within an array. The top left image in the figure shows the regularly spaced array of Figure 14. The remainder of the images show different random locations of the four element array and the resulting reconstructions.

The best solution is to allow the antennas to move periodically thus providing varying views of the image space and a variable aperture that also improves cross range resolution. In this type of scheme it is necessary to be able to determine the positions of the individual antennas in order for the image reconstruction to

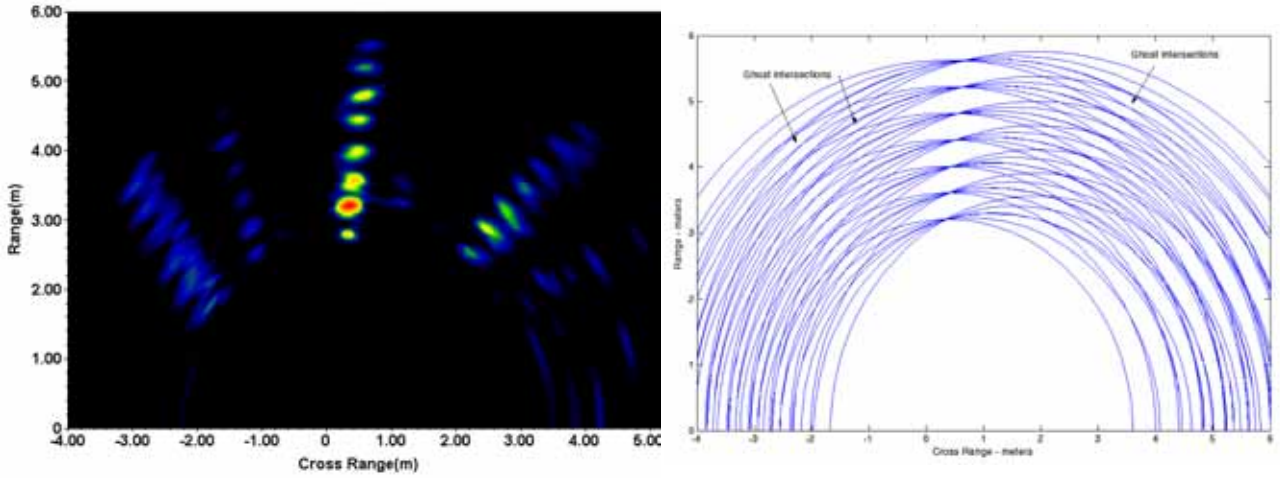

Figure 13 - Image ghosts are due to regular spacing of antennas and wall studs.

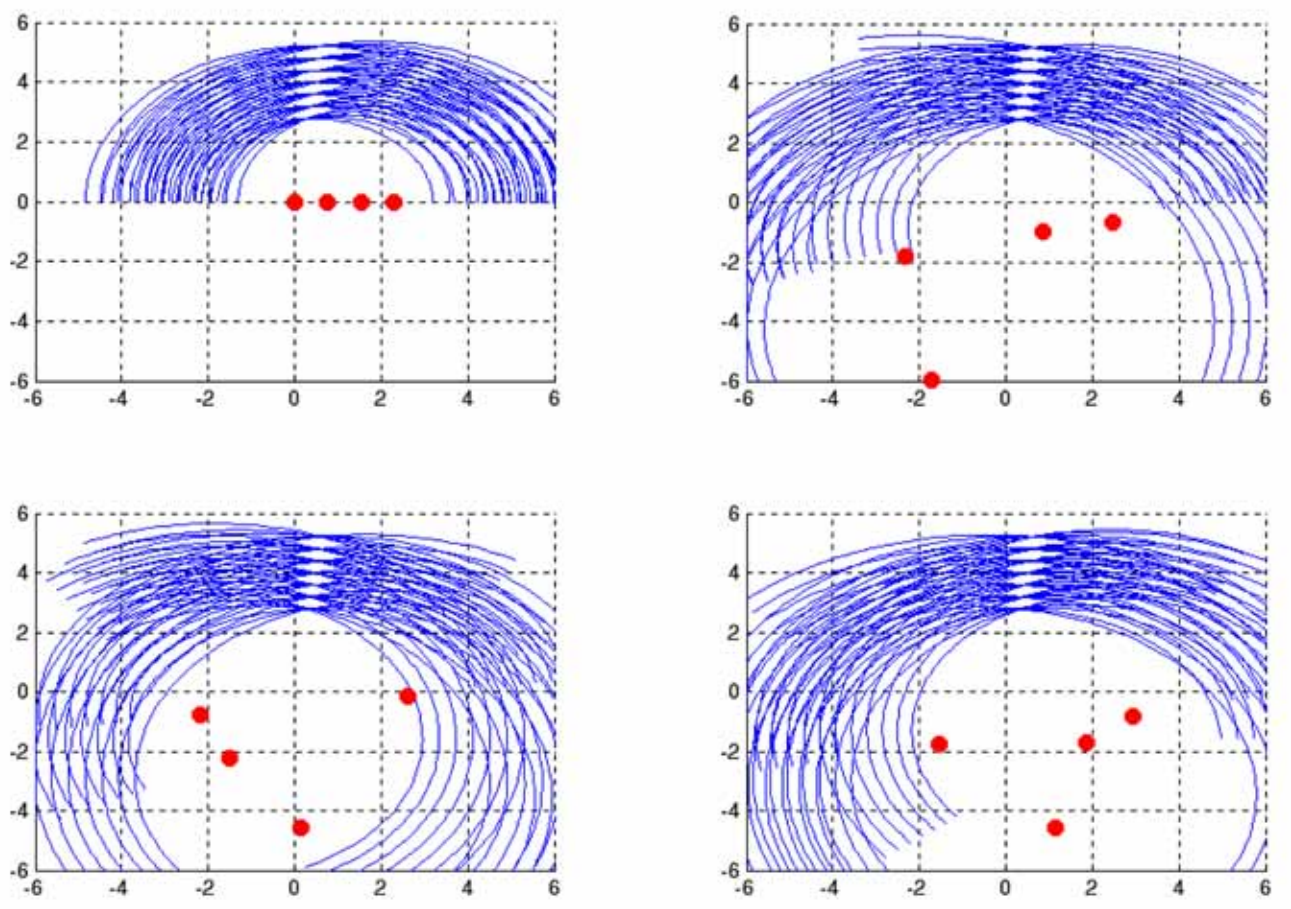

Figure 14 - Random placement of antennas reduces the effect of ghost images.

be successful. In a bistatic system, the direct path signal between the transmitter and receiver is the strongest and can be used to determine the distance between antennas.

Program Development Concept

Our program was designed to address the performance weaknesses just listed. The development concept is summarized in Figure 15. The major elements of the system are 1) small, portable radar sensors carried by individuals, deployed on vehicles, and/or set in place, 2) a processing and fusion network, and 3) command and display computer.

The radar uses direct digital synthesis. It is completely digitally controlled, hops at 10 μsec per frequency point, and can generate any frequency arbitrarily between 500 and 2000 MHz. This provides a low probability of intercept waveform, good jam immunity, user selectable range resolution, and a speed sufficient to generate coherent images suitable for using signal averaging to improve static imaging capability, and for detecting fast movement by coherent scene subtraction. Each radar sensor can both transmit and receive, but not both at the same time.

Concept Elements	Characteristics	Enables
Distributed sensor nodes on individuals on vehicles set in place	Frequency agile radar Coherent detection Random array Stationary or mobile	Dismounted operations Standoff surveillance Resolution where needed
Processing (fusion) network timing and control sensor location data and results	Imaging and motion 3 dimensions Weapons detection	Sensing through walls Sensing on the move Hostile ID
Command and Display individual mounted operation control command	Global views Local views	Situational awareness individual command elements

Figure 15 - Distributed, through-the-wall sensor network elements.

All of the sensor, processing, and display elements are connected through a wireless network. This network is used to provide timing and control to the radars to preserve coherent detection, to allow sensor movement, to transmit data to the processing units, and to receive processed image information. Using wireless network technology allows sensors to be added or removed from the network making this sensor useful for both large and small unit operations. If you view the moving sensors as moving antennas, it is easy to see that we remove the restrictions imposed by a linear, fixed array, allowing for different viewing geometries that reduce the effect of ghosts. With the sensors moving with individuals, the resolution of the imaging system stays best where it is most needed.

Processing is performed by specialized nodes that determine sensor location, fuse sensor data into coherent images, and perform other signal processing functions. As fusion points, these nodes have the information necessary to enable command elements to reconfigure the sensor network, generate selectable views of the operations space, and deal with lost or unavailable sensor nodes. For this program, there is just a single node that does the sensor control, processing, and display tasks.

The way the sensor works is very similar to the original brassboard system. The control node decides which sensor will transmit and broadcasts a message to the sensor network with that information. Those sensors that are not designated as the transmitter prepare themselves to receive. The control node then sends a message to the network that tells the transmitter to start a sweep.

In our brassboard system we had only a single transmitter and receiver so they had to be shared by all of the antennas. Consequently, we could only use a single pair of antennas at a time, stepping sequentially through all of the valid combinations in order to generate a complete image frame.

In our current system, each antenna has a radar sensor associated with it so all the sensors except the one transmitting can receive at the same time. This speeds up the process of acquiring a complete image frame over the method used by the brassboard, but puts more of a burden on the processing since the data rate is much higher.

The process is repeated until each of the sensors has been the transmitter, and then starts over again. A complete image frame consists of the coherent addition of the subimages from all of the combinations of transmitters and receivers - $(N_t) \times (N_r-1)$. As the number of sensors increases, it is possible to get good image results with only a subset of the sensors. For example, in an 8 sensor system after the first frequency sweep, there are 7 independent subimages that can be used to form an image. This image can be immediately displayed while the next sweep is being made. As additional sweeps are added to the image, the signal to noise ratio improves.

In order to maintain coherence, however, it is necessary for each receiver to have a phase reference from the transmitter. This is accomplished by using one of the sensors to transmit a beacon signal which is used to synchronize the internal processing clocks of the individual radars, using bidirectional scans to remove the residual uncertainty in the clocks in different sensor units, and then monitoring the drift in the high stability oscillators in each unit and correcting for the phase difference caused by the drift.

The control and processing node has a large computational burden put on it. Unlike the brassboard system with the fixed linear array, the physical location of each of the sensors in the networked system is not predetermined. However, in order to form images from the sensor data, it is essential that we know where each of the sensors is in relationship to one another. Fortunately, we have an advantage since a stepped frequency radar is also known as a high range resolution radar.

The largest signal at each receiver is the direct path signal from the transmitter. The Fourier transform of each frequency sweep generates a plot of amplitude vs. time. Using the time to the first peak in the transform and the speed of light allows us to calculate a distance from the transmitter to the receiver. Once a complete set of sweeps is made through all of the transmitter and receiver combinations, there is enough information to determine the position of each sensor relative to the others. In order to simplify the development effort on this program, the ability to form images while the individuals sensors were moving was not included. Sensor movement is still allowed, but image processing is suspended during movement.

The final element of the system is the wireless network. There is two way traffic over the network with commands sent to the sensors, and data returned from the sensors to the control and display node. This network communication is implemented with off the shelf commercial IEEE 802.11g wireless bridges manufactured by the DLink Corporation.

Hardware Development

Figure 16 shows the radar hardware we have developed. It consists of a set of circuit boards that are 4" in diameter assembled into a stack that, including space for a battery capable of sustaining operation for 4 hours, is about 8" tall and weighs 2.5 pounds with batteries and packaging. As shown in the figure, we have designed the radar so that it is modular, isolating the radar to a single board and providing power conditioning and control interfaces separately. The particular control interface that is shown was for development testing and wired applications. It receives commands and sends back data over a serial link at a rate of 4.096 Mbps. A special interface to this board has been developed to allow wireless operation with an IEEE 802.11g radio providing the necessary communication channel. Figure 17 shows the interface with a commercially available wireless bridge that provides wireless connectivity for the radar sensor.

Antennas are connected directly to the radar board through the connectors provided. For the independent radar sensor node configuration, a separate antenna interface board has been developed. It provides the connection to the radar board and acts as the ground plane for an inverted discone antenna as shown in Figure 18.

Antennas are connected directly to the radar board through the connectors provided, or for special applications, a separate antenna interface board can be developed. To support testing, we developed an interface board with an antenna switch that allows a single radar to scan an array of up to 8 antennas.

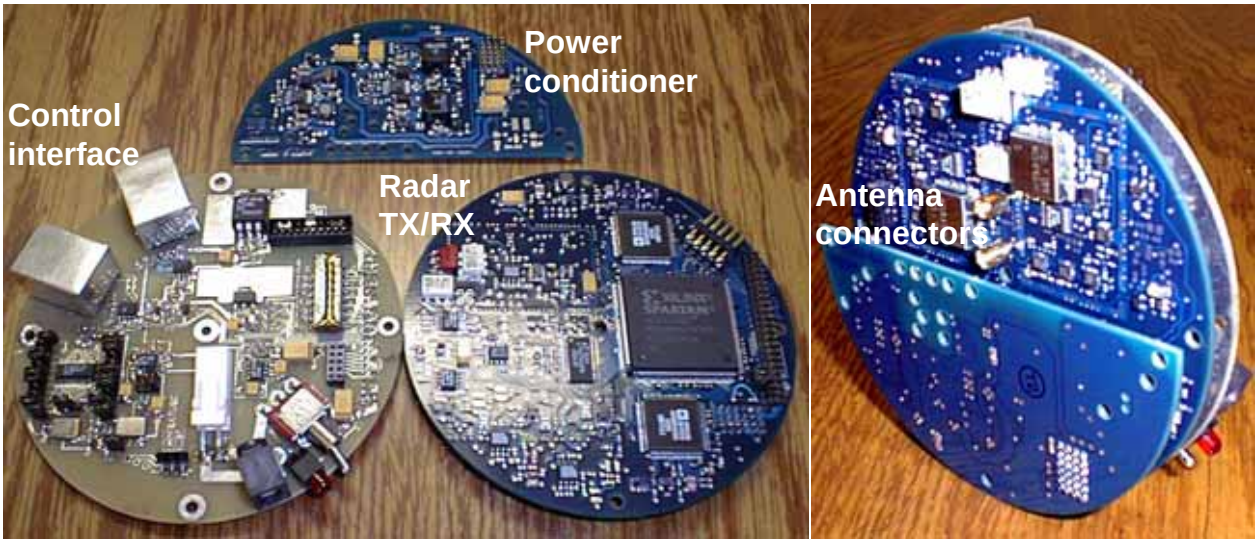

Figure 16 - Radar sensor node elements.

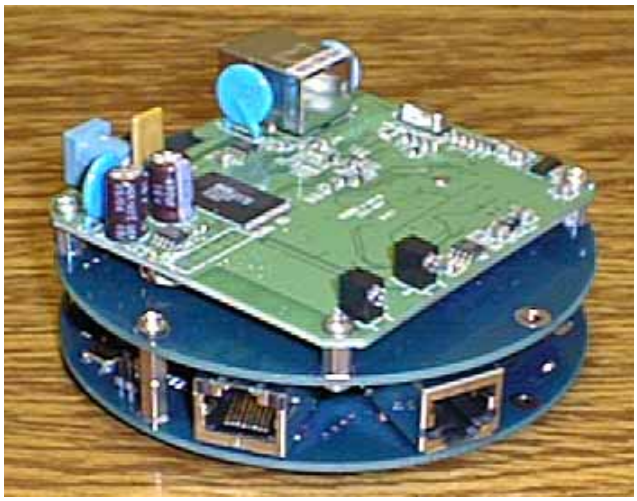

Figure 17 - Wireless sensor interface.

Figure 18 - Inverted dicsonce antenna interface.

Figure 19 shows a simplified block diagram for the transmitter portion of the radar. The transmitter uses a direct digital synthesizer (DDS) as a reference frequency generator that is used to control the output of a voltage controlled oscillator (VCO) operating over the range of 1 to 2 GHz. The output of the VCO is divided and fed into a phase/frequency detector along with the output of the DDS. This phase locks the VCO to the DDS. For frequencies between 500 and 1000 MHz, the VCO output is divided by 2. By using a DDS we have precise control over transmitter frequency and can arbitrarily select frequencies. A field programmable gate array (FPGA) is used to decode commands from the radar interface board to set up the DDS to generate a frequency scan.

The receive section of the transmit/receive board mirrors the transmit design on the front end, using a DDS to generate a frequency that is offset by 10.7 MHz from the transmit frequency, fed into a mixer with the received signal, and filtered to produce the intermediate frequency (IF). However, to obtain the design goal sensitivity for the receiver, it was necessary to eliminate a second down conversion to a lower IF as is common in most designs. A digital demodulation technique going directly from the first IF to baseband is used. In this technique, the IF is digitized at four times the IF and the samples are manipulated to reconstruct the inphase and quadrature portions of the received signal. This reconstruction is performed in the FPGA.

With both the frequency generation and demodulation portions of the radar being controlled by the FPGA, the radar has become what is commonly called “software defined”. All of its essential operating parameters are controlled by software giving it the ability to be reconfigured to meet operational requirements. This

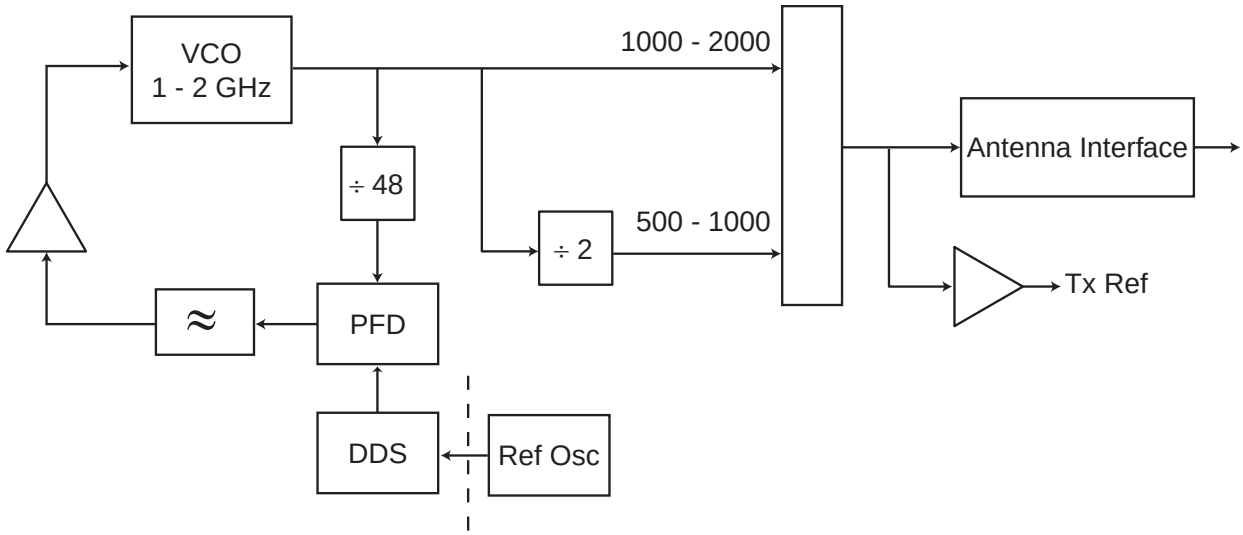

Figure 19 - Simplified transmitter block diagram.

approach has allowed us to simplify the hardware portion of the radar but at the cost of more complex firmware that must be developed for the FPGA.

The heaviest element of the radar system is currently the battery. It is a 48 Watt Lithium Ion Polymer rechargeable pack that was sized to provide 4 hours of continuous operation. The radar currently consumes 13 Watts while sweeping, 8 Watts in standby, and 3 Watts in sleep mode with the largest power consumption being the DDS chips and the high speed A/D. A second generation of lower power DDS chips has recently been introduced and new A/D converters have appeared that are more power efficient. Incorporation of these new electronic elements would reduce total radar power consumption to about 5 Watts saving additional size and weight associated with the battery.

Software Development

Three different types of software programs were developed during the program to support hardware checkout and debugging, algorithm development, and field testing. The first program was designed as a standalone program that would allow us to check the communication protocol that had been set up for controlling the radar operating parameters and allow us to use it during the hardware checkout and debugging process. The DDS operation is controlled by setting values in a series of registers that determine what frequency is generated, the type of frequency sweep, whether the DDS is powered up or down, the sweep rate, and the status of control lines that are used for IF gain and antenna settings.

The program was developed as a windowed interface allowing the user to select preset commands and values, build commands for debugging purposes, view the contents of the DDS registers, and receive data

from the FPGA that arbitrates the communication between the PC and the radar board. Figure 20 shows the program user interface.

This program uses a set of TCP/IP software that implements the UDP protocol. This allows us to use either a wired or wireless Ethernet interface to the radar. As shown in the figure, the radar being tested is identified by an IP address. Connectivity is first determined by sending a command using the Quit button to which the radar responds with a question mark and carriage return. The two large windows to the right in the interface show the sent and received command sequences in both ASCII and Hex. To the left side of the interface are a set of fill in boxes and buttons that allow individual sweep parameters to be sent to the radar

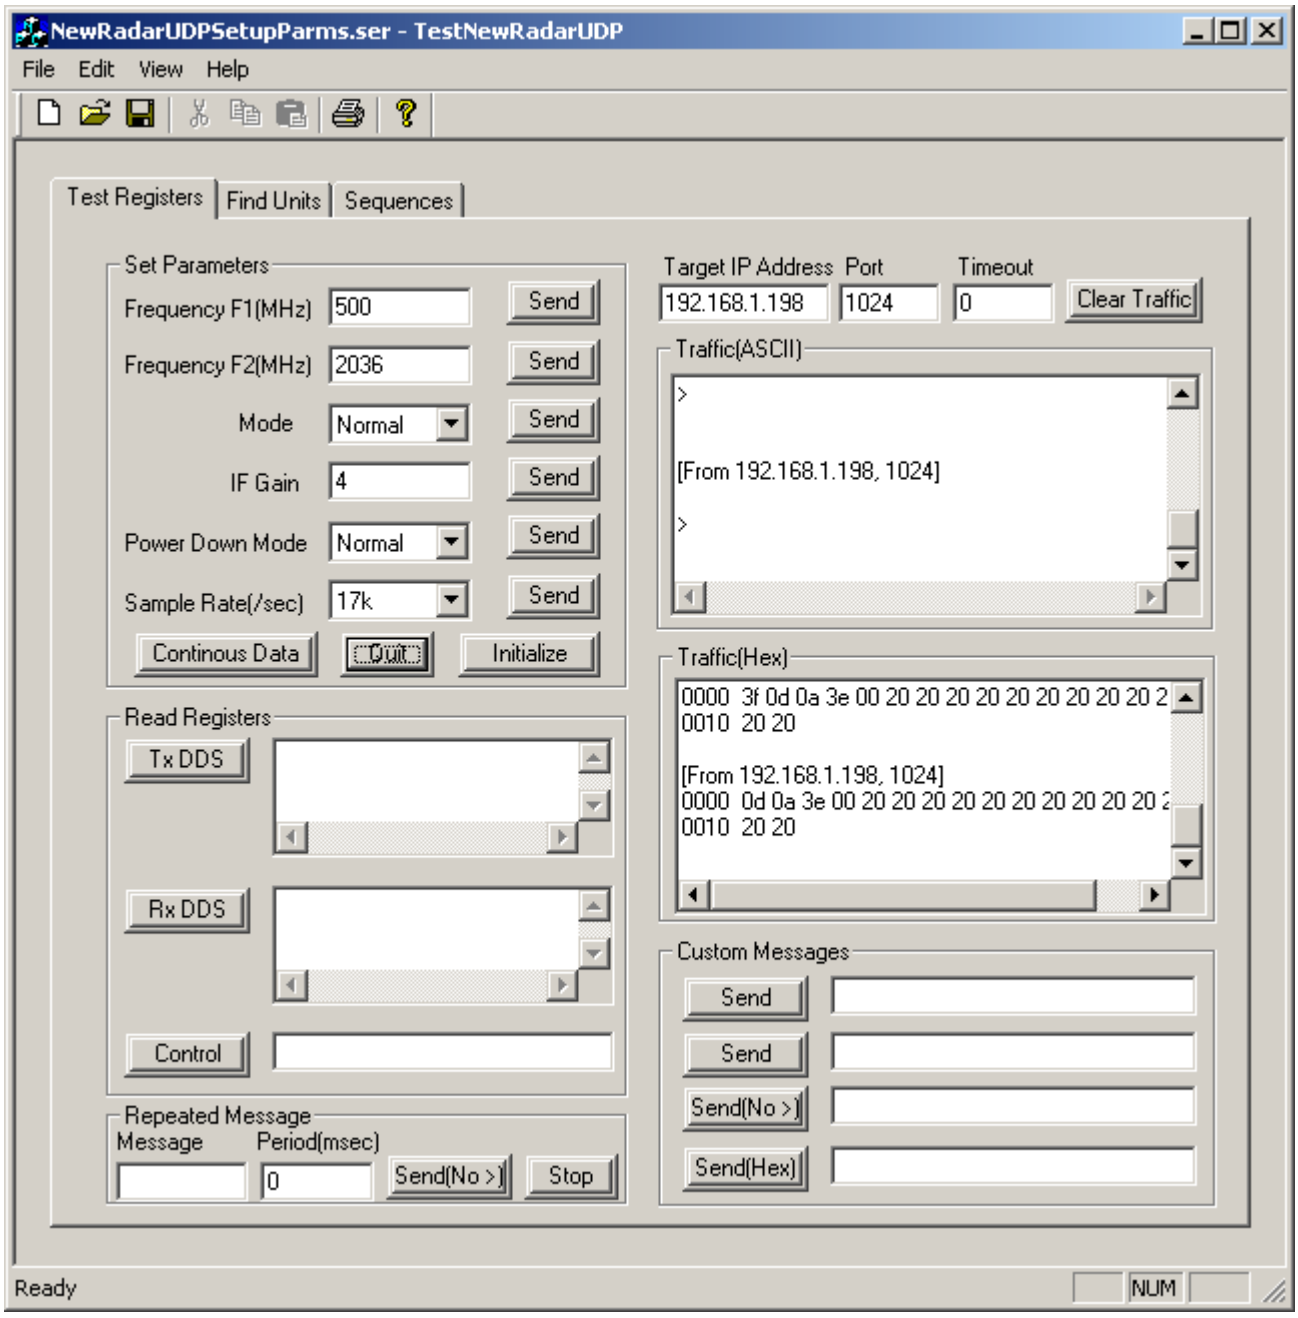

Figure 20 - Standalone radar test program user interface.

and the various registers of the radar to be read to determine whether the sweep parameters were successfully implemented. To the bottom right of the interface are a set of fill in boxes that can be used to send individual commands to the radar. This interface was used extensively during system development to identify the source of unexpected radar behavior, assist in finding radar setup problems with the other software programs developed to perform data collection and imaging, and to identify the source of communication problems in the sensor network.

The program used for all of the control and processing functions was based on the one originally developed for our precursor brassboard imaging system. All of the test development software created for the standalone program was incorporated into a copy of the program and used as a basis for further development. The new version of software implements the command structure for the radar and adds the ability to have antennas in arbitrary locations.

Our control program is implemented using Visual C++ and the Microsoft Foundation Class (MFC) library. The Multiple Document Interface(MDI) enables the simultaneous display of results (plots, tables, etc.) from more than one data source (files, data streams, user input, etc.) and employs a Document/View architecture. A Document is an object that is typically responsible for collecting, manipulating, and storing data. Each Document object type may have one or more View object types associated with it, and these objects may be displayed as the program sees fit. Each View object embodies a visual presentation (plots, tables, forms, etc.) of some or all of the data from the associated Document object, i.e. it typically controls a window whose contents are a representation of the data.

Figure 21 shows the resulting user interface. The major functions of the program are accessed from the File, Edit, View, and Tools menus at the top of the main window. These functions enable acquisition of single image frames or sets of image frames from both single and multiple unit configurations of the radar system. Sets of image frames are used for motion detection studies. Single frames are used to help define low level signal processing parameters that are applied to the image and motion detection algorithms. When the desired function is selected from one of the menus, the program launches a form that allows the user to specify the appropriate control parameters and then start the selected task.

System and processing configuration information is entered in the Analysis Parameters form shown in Figure 22. It has eight tabs that allow the user to specify antenna location and cable delays, the area over which an image is reconstructed, the number of FFT points to use for signal analysis, files to use for background subtraction and signal normalization, and the values to use for various correction factors that account for special hardware configurations. It is used primarily for single unit parameter configuration.

Other parameters that may be selected are the number of sweeps to average while taking data (to build up a higher signal to noise ratio data set), whether to display the data graphically as it is being collected, the name

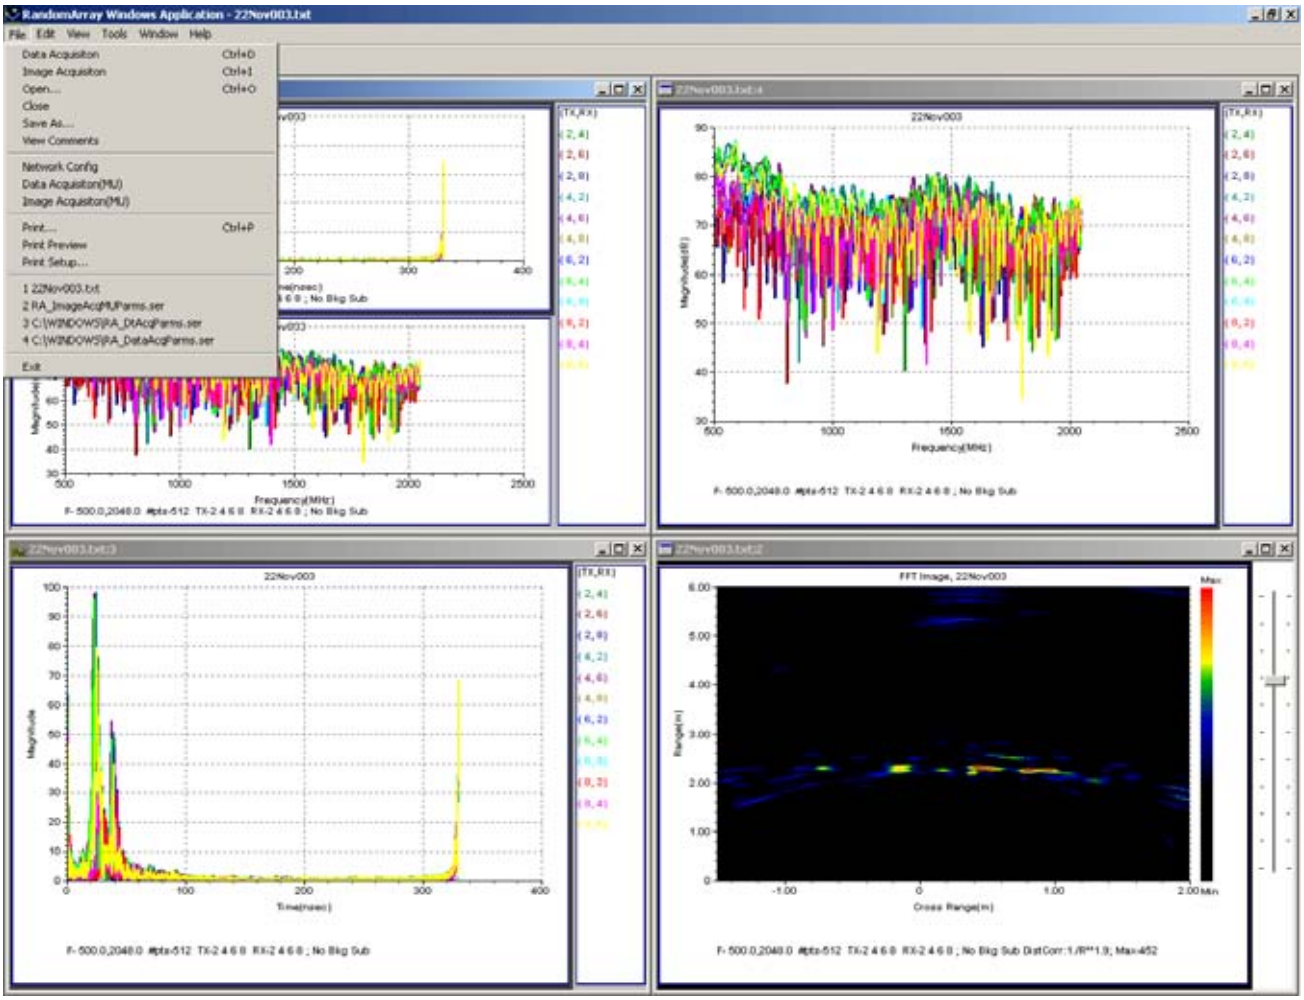

Figure 21 - Radar control and analysis program user interface.

to give the resulting data file, the directory path where the data file is to be stored, whether to automatically increment the file name for subsequent tests, and which IP address to use for communication. Maximum sweep speed (samples per second) is determined by the settling time of the hardware frequency generation circuit within the radar.

Figure 23 shows the single unit data acquisition function selection form. This form is where the operating parameters of the radar are specified before the beginning of any test series. The user can select the start and stop frequencies of each sweep, the number of frequency points for each sweep, and the sample rate per frequency point. While the user may select arbitrary frequency ranges and numbers of points, the physical limits of the radar impose some restrictions. Because of memory limitations, the radar can not support frequency sweeps with more than 4096 points. Perhaps the most significant part of this form is the portion of the radar settings tab that allow the user to exclude frequency ranges during a frequency sweep. This behavior is possible because all of the radar operating parameters are software defined. It is important because users are often worried about the interference that wideband systems create when passing through communication, navigation, and other important operational frequency bands. The frequency exclusion

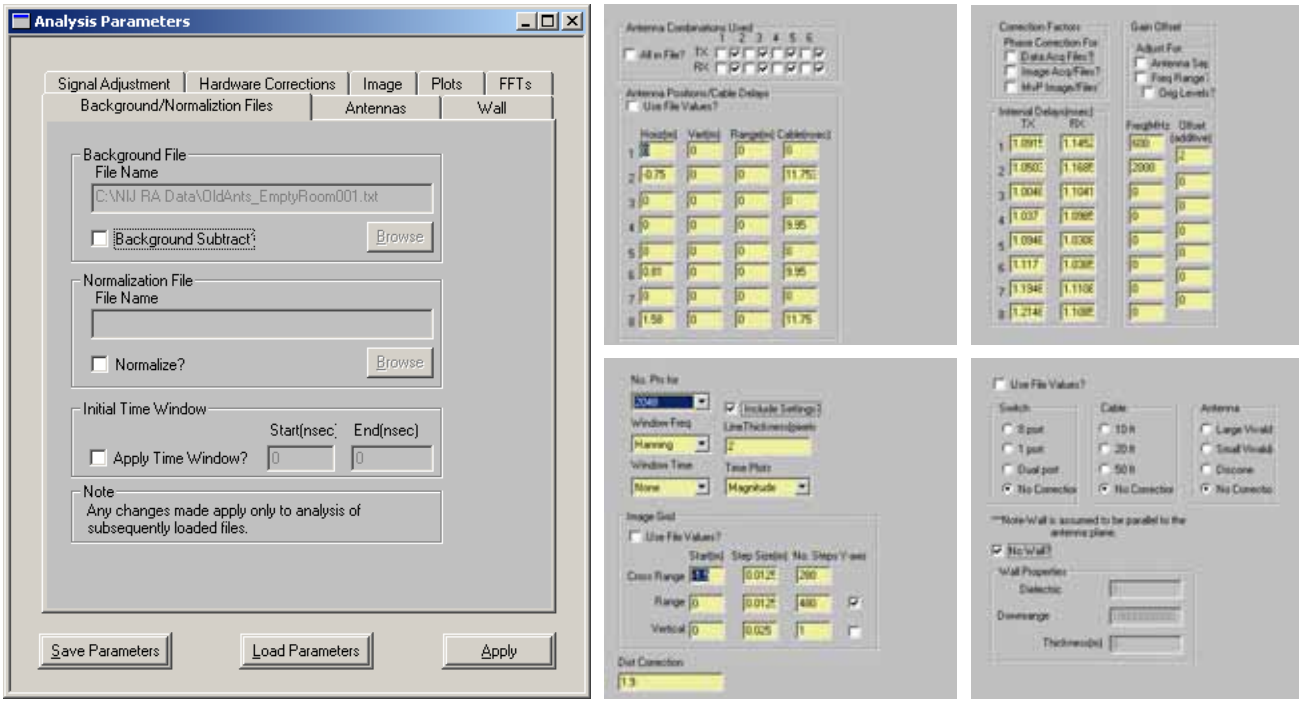

Figure 22 - Single unit analysis parameters form.

bands shown in the figure are typical of those that military frequency spectrum managers require omitting when granting permission to operate on military installations. This generally causes problems for wideband systems that do not have the capability to selectively control emissions. Figure 24 shows an example frequency sweep where these bands have been excluded. Omission of selected frequencies does not

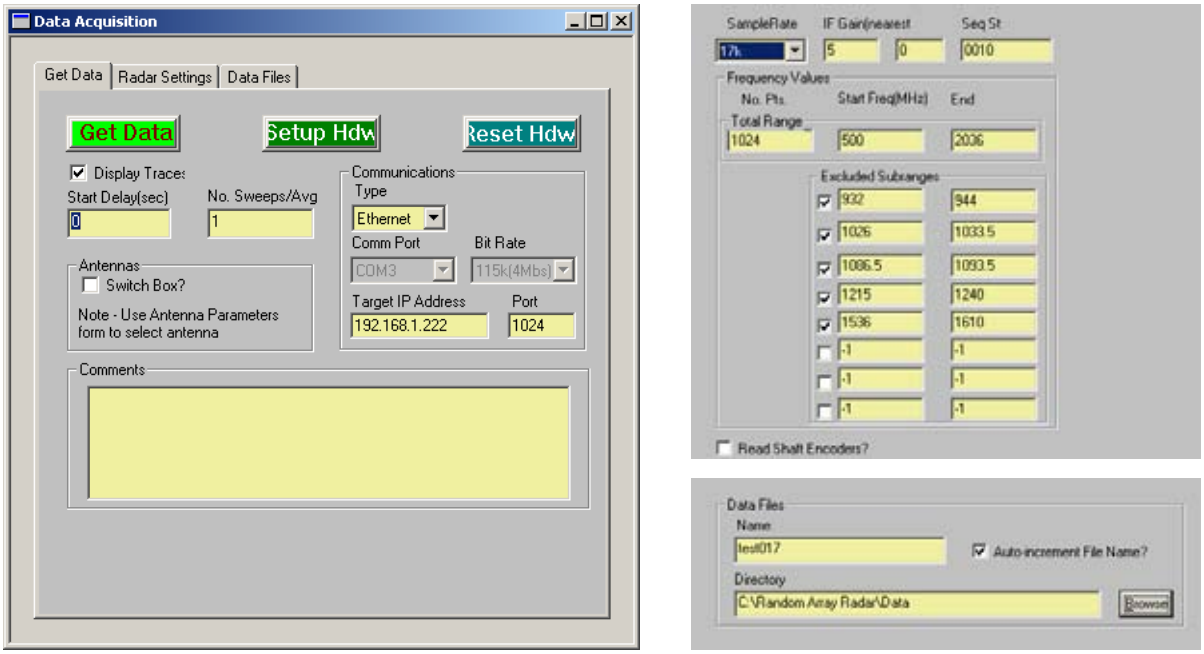

Figure 23 - Single unit data acquisition form.

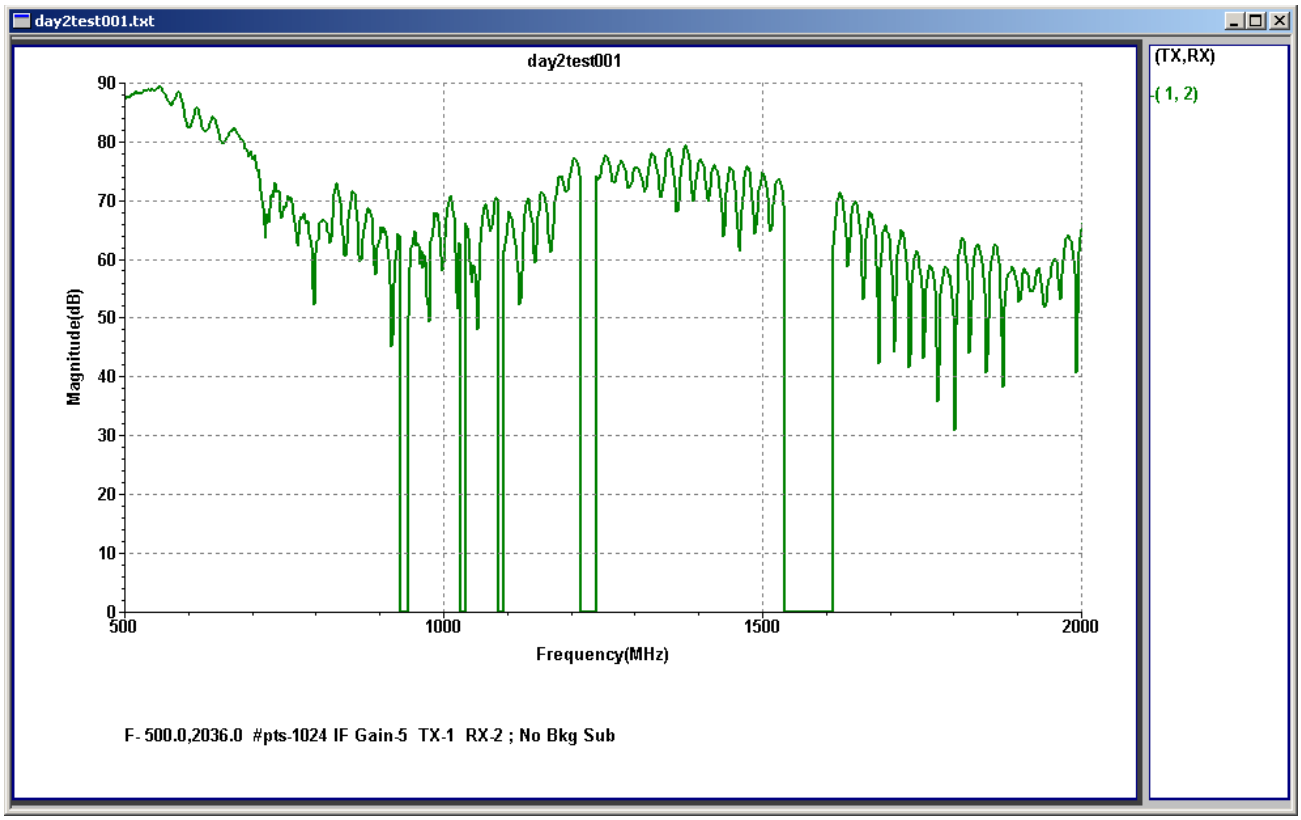

Figure 24 - Selective frequency omission during radar operation.

change the range resolution which is set by the bandwidth. It reduces the signal to noise ratio since there are fewer data points taken over the same time span.

Image acquisition for single radar units is controlled by the form shown in Figure 25. As with the data acquisition form, the user controls all radar sweep parameters through this form. In addition, there are user defined parameters that control the area over which the image is reconstructed, distance corrections that can be made independently to the static background image and the moving target image, the type of motion detection algorithm that is used for displaying the data, and other signal processing operations that can be used to manipulate the image data. Since data is automatically saved each time the system is used, this form allows the user to recall previously saved image files to play through the software and to explore the effects of the different signal processing operations on image results.

As is the case for single units, there is a set of forms that allow the user to set up the radar sensor network and configure it to collect and produce images from single or multiple frames of data. Figure 26 shows the form for multiple unit parameter configuration. It has eleven tabs that allow the user to specify radar operating frequency range, number of points, gain, frequency exclusion ranges, and specific sensor physical configuration information. An important part of this form is the network configuration information. We have allowed for mixed sensor combinations to allow both fixed and mobile sensors to be used together for forming images. Fixed sensors may either be standalone single units, or single units with a multiple antenna

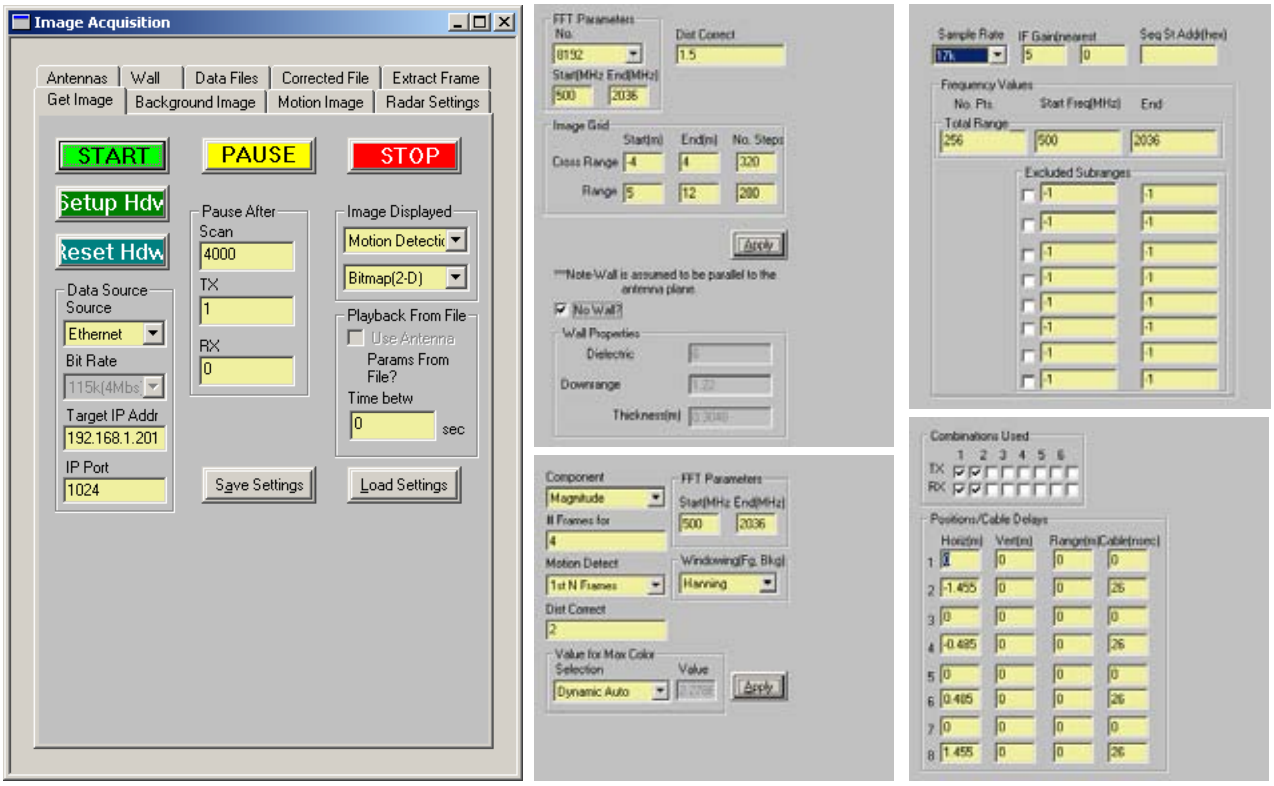

Figure 25 - Single unit image acquisition form.

switch. The tab settings allow for 4 fixed units with multiple antennas for which all of the antenna position, cable delay, and antenna switch delays are specified. For cases where there are no multiple antenna units, the user can enter setup information for up to seven individual units in the multi unit tab forms.

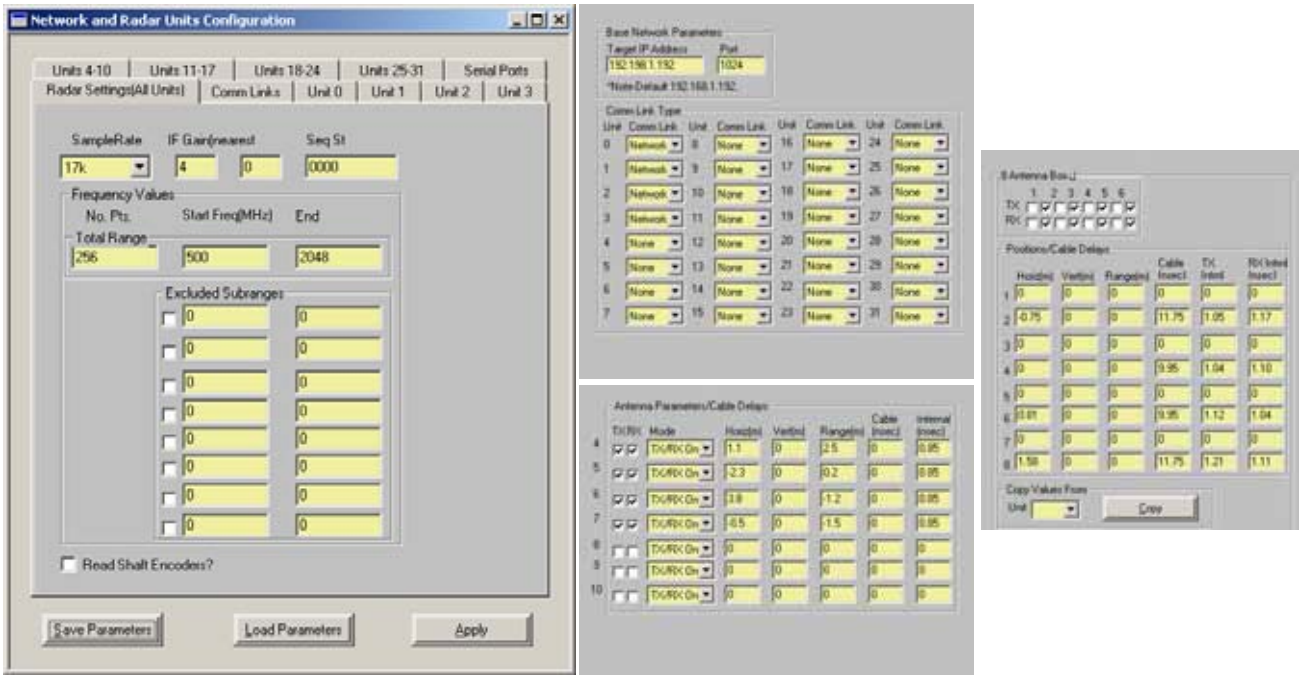

Figure 26 - Multi unit parameter configuration.

The communication links tab allows the user to specify the network address of each sensor in the system. We have assigned all radar hardware units individual IP addresses with the base address starting at 192.168.1.192. In this scheme unit zero has the base IP address while unit 5 would have an address of 192.168.1.197.

As shown in Figures 27 and 28, the multiple unit data and image acquisition forms are very similar to those for the single unit. The major exception is the millisecond and oscillator synchronization buttons. These buttons give the control software the ability to synchronize the independent radar unit clocks so that they can collect data, and to periodically adjust the phase of the clocks to correct for drift of the independent clock oscillators in each unit.

Because the radars are operated in a bistatic mode where one transmits and the others receive, it is necessary that each radar be on the same frequency at the same time in order for successful operation to take place. At the highest hopping rate of the radar, each frequency is transmitted for only 11 μseconds so the time window available for reception is small. To solve this problem we implemented a technique where a single frequency beacon signal is transmitted from a single radar unit to all of the others and then is turned off. Each receive unit monitors the level of the beacon signal and when it falls to half of its maximum average value, sets its internal program clock. In that way, we are ensured that each radar is executing the same instruction at the same time within the uncertainty of the clock synchronization. Figure 29 shows a set of

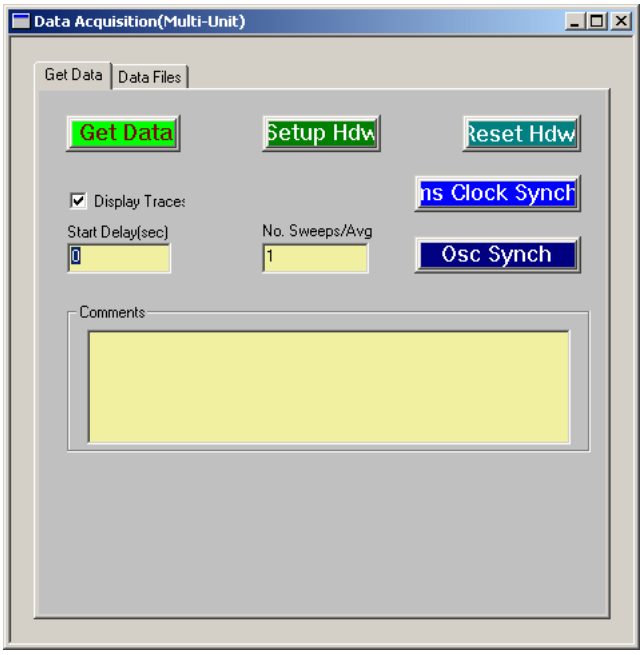

Figure 27 - Multi unit data acquisition form.

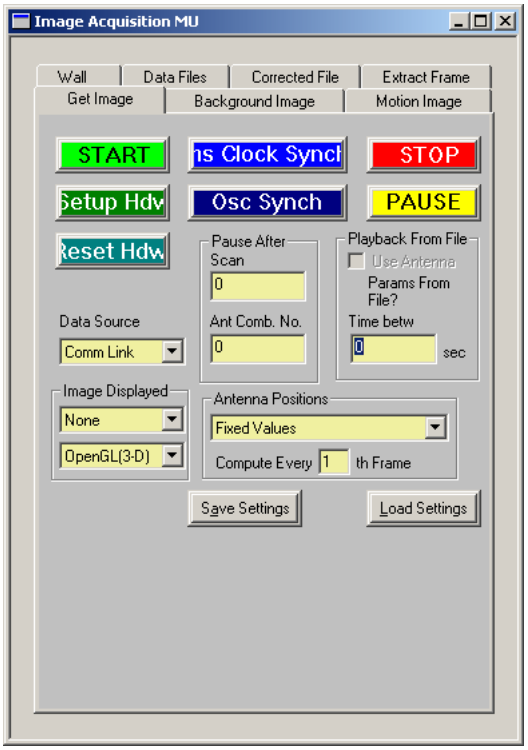

Figure 28 - Multi unit data acquisition form.

oscilloscope traces that illustrate this process. Because there are physical delays in the radar hardware, there is a residual uncertainty in the start times of the individual clocks of plus or minus 60 nsec. The millisecond sync button on the multi unit data and image acquisition forms implements the commands necessary to align these internal clocks.

Oscillator synchronization is also important. There are independent clock oscillators in each of the radar units which are used as a reference for all timing functions in the radar units. They are high stability clock oscillators manufactured by MTI Milliren. All oscillators will drift over time with differences between units in not only rate of drift, but also direction. The frequency drift differences between units cause phase differences that accumulate over time and ultimately destroy the phase coherence between the radar units and, therefore, their ability to perform imaging. The effect of this drift is shown in Figure 30 for a drift rate of 1 part in 10^9. The first panel of the figure shows the range profile for a set of radars with no phase differences in their clocks. The remaining three panels show the result of clock drift after 0.1, 1, and 10 seconds

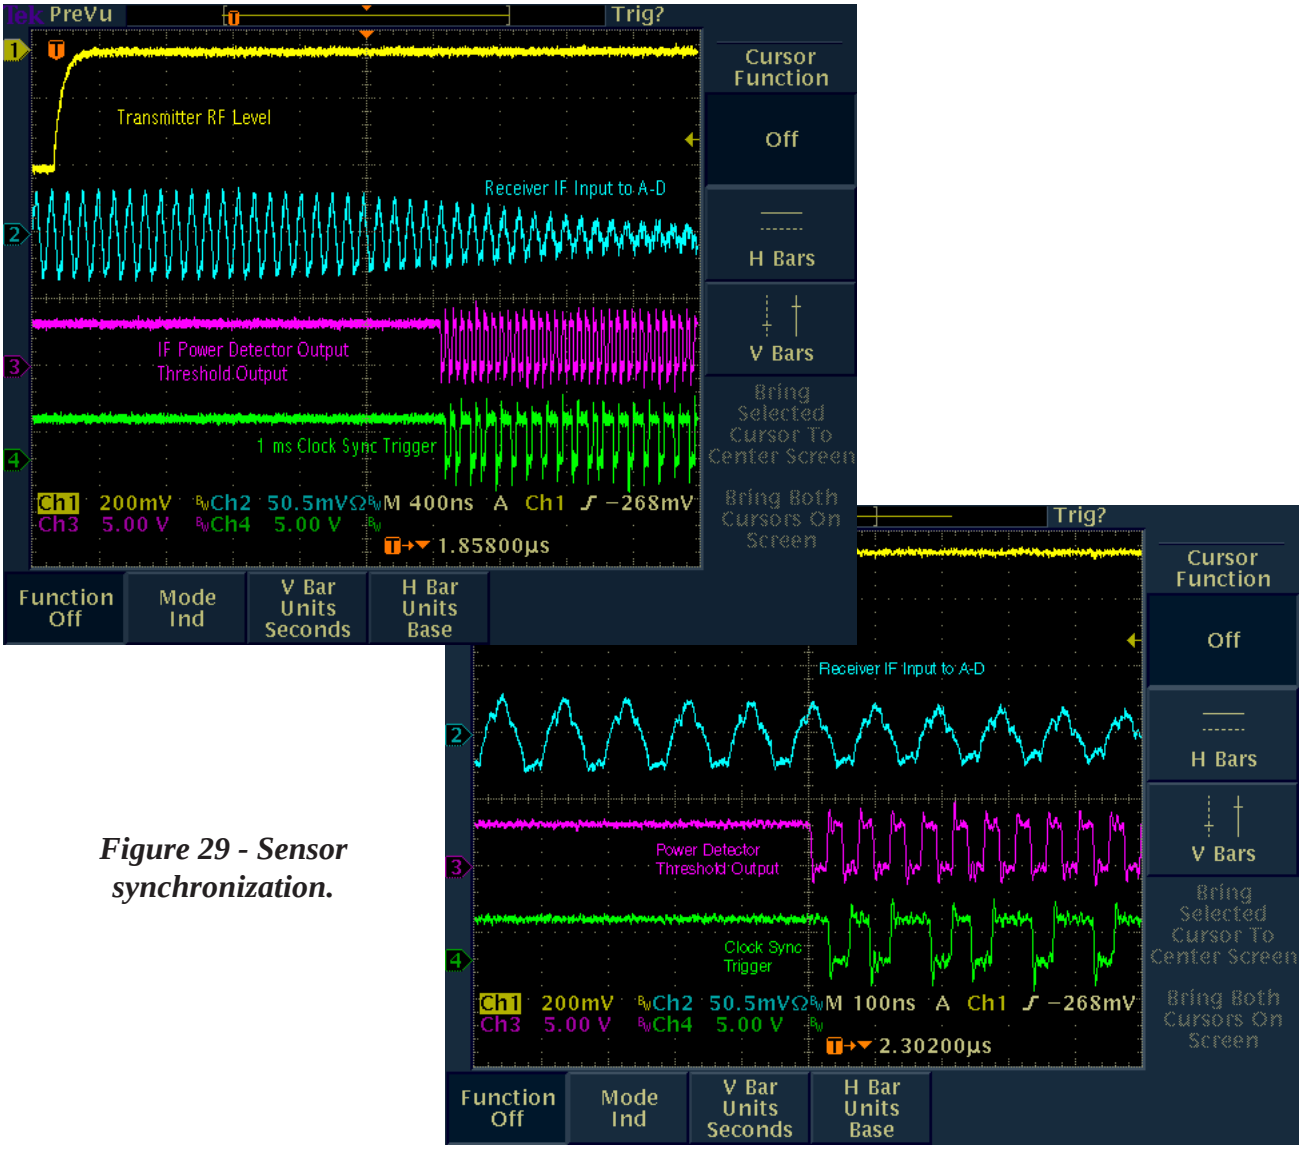

Figure 29 - Sensor synchronization.

respectively. As seen in the last panel of the figure, as the phase error accumulates, the sharpness of the peak is destroyed, with the result that the radars will ultimately lose lock with one another. By monitoring the drift of the radar units and periodically correcting it, the imaging capability of the sensor system is preserved. The oscillator sync button in the control software interface implements this correction process.

The third software program that was developed on this program was designed to explore the feasibility of using modern graphics processors that are designed for video gaming to render the images for our system. We succeeded in creating an interface that collects data from the radar, performs all of the calculations needed to create an image, and then displays it using the graphics processor. This process allows the user to move arbitrarily around the image while the data is being displayed, to interactively change the frequency band of the data being displayed, change the threshold for the image color map being displayed, and move the portion of the image being displayed. In addition to pan, tilt, and zoom functions, we created a fly through visualization mode that has an identical mouse and keyboard interface to the currently popular first person shooter video games. With this interface we have been able to process and display images at 10 frames per second on a 1.4 GHz desktop computer. It has been used extensively during our experimental testing to view and render image data in real time and has the ability to change the type of motion detection algorithm used. Figure 31 shows the user interface from this program that has been implemented using Open GL graphics.

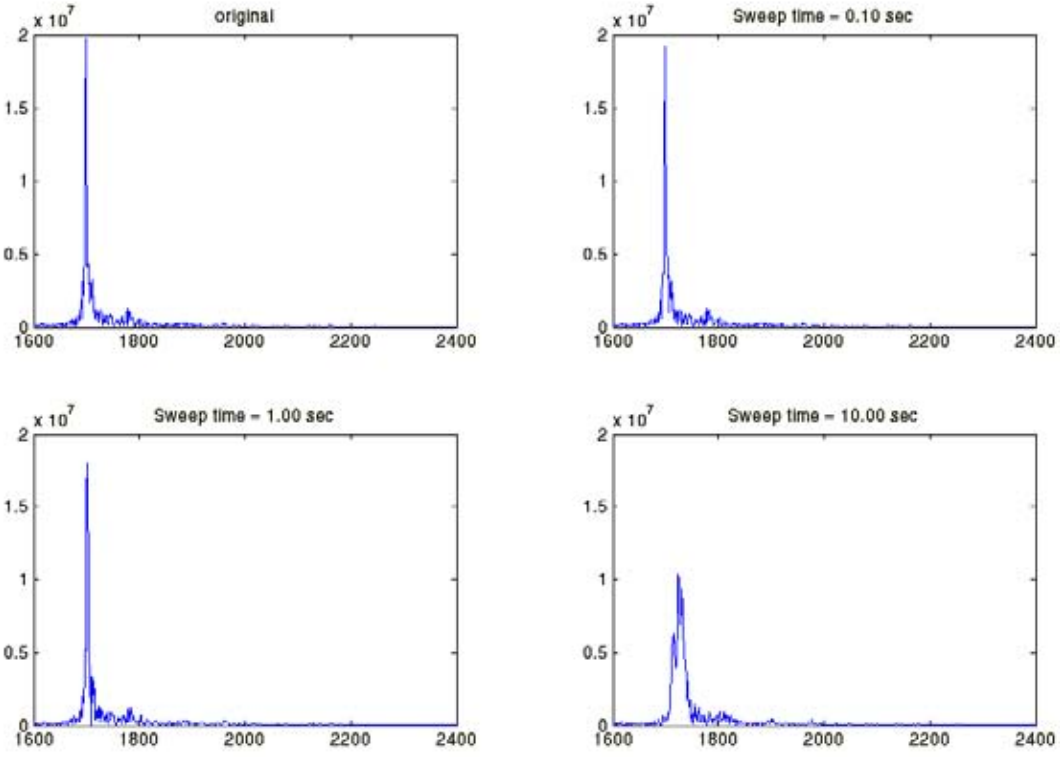

Figure 30 - Effect of oscillator clock drift of 1 part in 10^9.

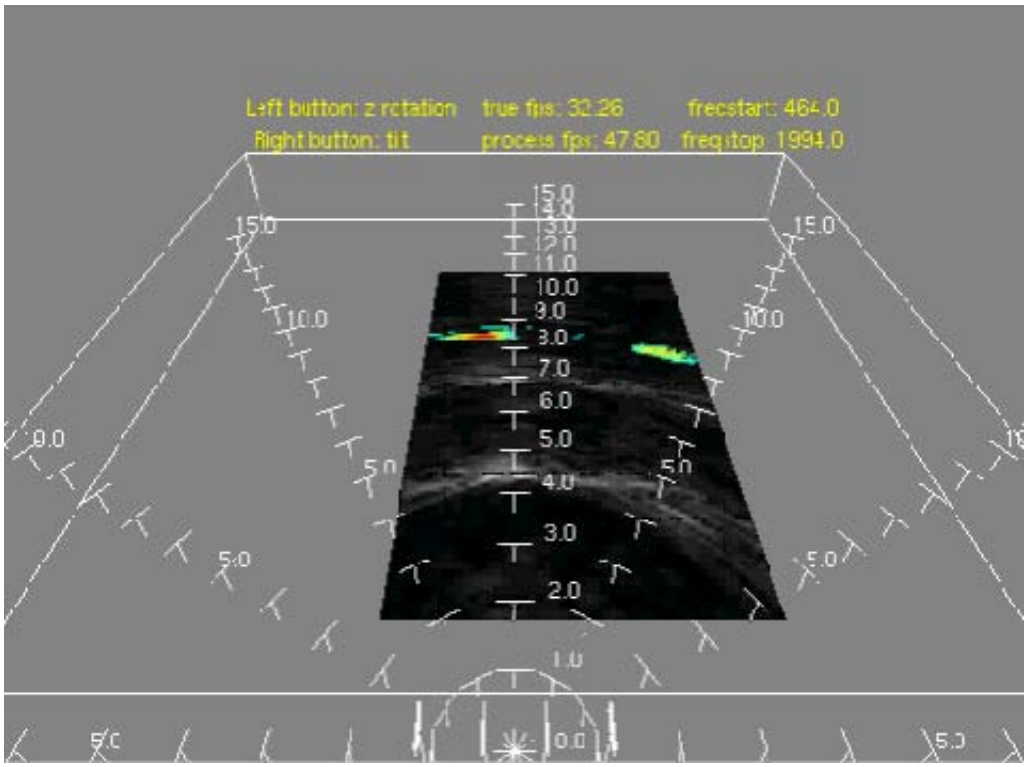

Figure 31 - Imaging using Open GL graphics.

Experimental Program

Extensive data collection was performed on this program. In order to get data early in the program with which to verify radar operation and enable algorithm development, a special antenna interface board was constructed. This interface board, shown in Figure 32, enabled us to use a single radar with an array of up to eight antennas. In this configuration, the antennas were used a single pair at a time with the data from all pairs combined to form a single image frame. To make testing at different locations easy to perform, we constructed a collapsible array with four antennas. This configuration was similar to the original brassboard imaging array of our precursor program, and enabled us to gather similar data for performance comparison. Figure 33 shows the array in the collapsed position, and Figure 34 in the open position.

To take advantage of the radar's capability to sweep frequencies at a high rate, a communication interface board with a high speed RS-422 serial port was developed and used with a commercially available RS-422 serial board on the PC side. This allowed the radar to send data at 4.096 Mbps and permitted the collection of data at a rate sufficient to perform motion detection in real time. Figure 35 shows the setup for testing the motion detection algorithms and a static image that was extracted from the data. As shown in the figure, the radar is at zero range with an 8' x 8' stud wall faced with a single sheet of drywall at a distance of 2.5 meters (8.2') from the radar. In the test, a person walked from behind the radar around the left side of the wall, circled around the chair at 5.8 meters (19') and then sat down attempting to stay as motionless as possible. After a period of time, the person got up out of the chair and walked back toward the radar and

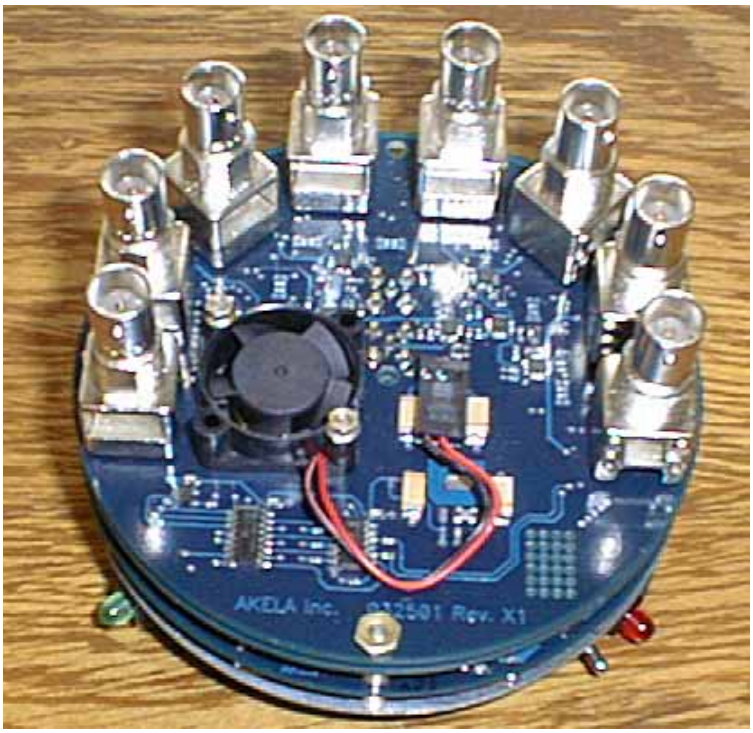

Figure 32 - Eight antenna test configuration.

Figure 33 - Collapsible array.

Figure 34 - Array in deployed configuration.

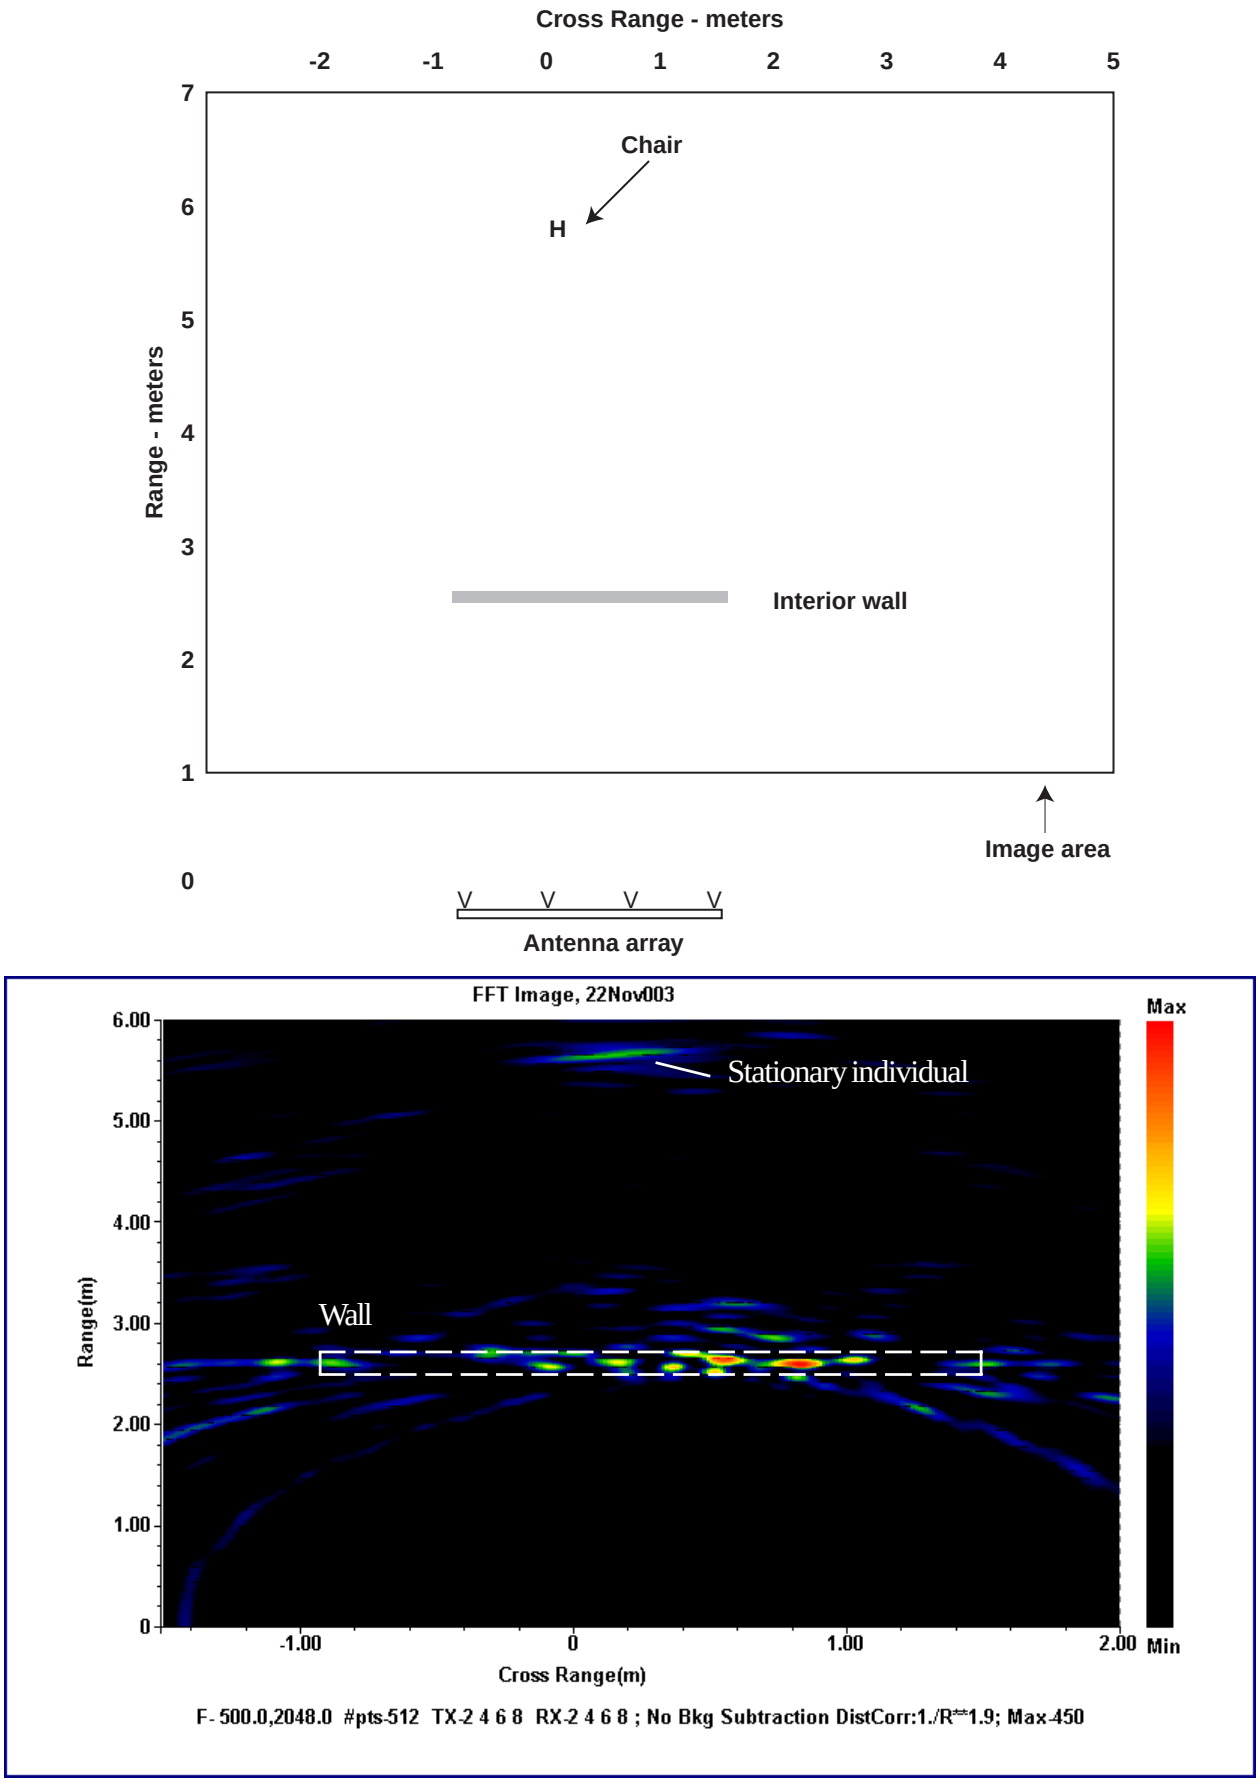

Figure 35 - Test geometry for breathing detection and static image.

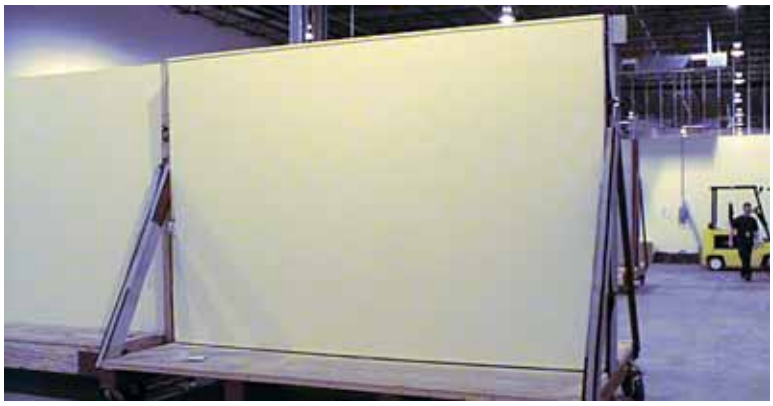

Drywall

Brick

Concrete block

Adobe

Figure 36 - Test walls.

Reinforced concrete

Figure 36 - Continued.

around the left side of the wall. The person was detected while sitting motionless, and this was the first demonstration that our radar was sensitive enough to detect breathing.

The first extensive set of testing with the radar was performed in cooperation with the Army at the I2WD wall facility at Fort Monmouth, NJ. This test facility has five different walls, as shown in Figure 36, made of drywall, brick, unfilled concrete block, adobe, and steel reinforced concrete. Each wall is 12" thick, 8' high, and 10' wide with the exception of the steel reinforced concrete wall which is only 8' wide to stay within floor loading limitations of the test facility. The drywall wall is hollow inside and faced on both sides with sheetrock. The brick wall has a brick facing on one side and drywall on the other.

Testing was conducted over a period of two days to test the high speed data collection capabilities of the radar, and to gather a set of data to use for algorithm development and for determining the electromagnetic material properties of relevant wall types. We used two different radar configurations with wired connections to a set of antennas. Figure 37 shows the field tests radar hardware used. In each enclosure is a command interface board, power supply board, radar board, and antenna switch board - one with provision for eight antennas and the other provision for two. Each radar enclosure contains a battery capable of operating the radar continuously for 4 hours and provides a high speed serial interface that allows a control computer to configure radar operating parameters and receive radar data. Since the eight antenna switching board reduces the transmitted signal, especially at the higher frequencies of operation, we used the two antenna configuration for static testing where we desired higher signal to noise ratio data.

Both the wideband attenuation and the frequency dependence of the dielectric constant of each wall material was measured using the same test setup. To measure the attenuation, two antennas were used in a bistatic configuration, separated by a fixed distance but pointed directly at one another without a wall present. In this configuration, the first signal received is the direct path between the two antennas. Applying a narrow

Figure 37 - Radar field test enclosures.

window to this direct path signal and then Fourier transforming to the frequency domain, gives the frequency dependent response of the radar system. Performing the same measurement, but with a wall between the two antennas, gives a result that is the response of the radar convolved with the response of the wall. Using the same time window as on the clear path and transforming to the frequency domain, then dividing the wall plus radar data by the radar alone data, removes the effect of the radar system response and leaves the normalized frequency dependent attenuation response of the wall.

The bulk dielectric constant was measured with the same configuration. In this case, it was necessary to measure the distance between the two antennas and the thickness of the wall. Because the direct path between the antennas is the first and strongest signal, comparing the times at which the peaks in the radar only and radar plus wall data occur is a direct measure of the time delay introduced by the presence of the wall. By knowing the actual distance between the antennas, the thickness of the wall, and the time difference between the peak response, it is possible to calculate the path length difference and, therefore, the bulk dielectric constant in the wall material. The last photo of Figure 36 shows a photograph of the test setup for measuring the material properties of the concrete wall.

We took ten sets of data for each wall for a variety of radar, antenna placement, and software correction conditions. Our control and analysis software had been previously modified to allow us to select corrections for different type of antennas, cable lengths, and radar antenna switch configuration so we could perform preliminary analysis at the test site and adjust our testing plan accordingly. For each wall we ended up with a consistent set of measurements with antennas placed right at the wall, and at 30”, 48”, and 72” away from the wall.

All motion detection data collection was performed with the eight antenna switch configuration and an array of four antennas. We performed three motion tests on the adobe wall, seven on the concrete wall, and two

each on the block, brick, and drywall walls. The antenna separation and orientation of the antenna array with respect to the walls was limited by the lengths of our antenna and power cables. This resulted in a range of standoff distances for the motion data of between 4 and 20 feet. In addition, we performed two long range tests on the concrete wall where the antenna array was 40 feet from the wall. For each wall we collected data of a person walking around behind the wall for testing the large motion algorithms, and of a person sitting quietly in a chair for testing the small motion algorithms.

Since the test walls were mounted on elevated platforms, we placed radar absorber both in front of and behind the platform base in order to reduce the chance that radar energy would bounce underneath the platform and give us a false indication of movement. Figure 38 shows a typical breathing test setup. The flat sheets of radar absorber for preventing multipath from the floor of the test facility can be seen clearly on the front side of the wall.

The laptop computer used for testing had a special high speed serial interface for communicating with the radar. Because of limitations in the operation of the special interface, our motion testing was limited to a sample rate of 30,000 samples per second - about 1/3 the maximum radar rate. This caused the resulting movement data to be more coarse than with the highest radar sample rate, but increased the signal to noise ratio of the data.

Attenuation properties of the walls was measured directly from the wide band frequency response of the wall. Because the radar has a frequency dependent response, it is necessary to remove it from the data in

Figure 38 - Test setup for collecting small motion data.

order to get the true frequency response of the wall. Figure 39 illustrates this dependence graphically. The graph on the left shows curves for the through transmission response of drywall and for the antenna to antenna response of the radar with no intervening wall. As can be seen in the figure the responses are nearly identical since drywall has very little attenuation. The graph on the right shows a similar set of curves for the reinforced concrete wall. In this case the effect of attenuation through the concrete is evident in the lower magnitude of its response curve, however, its overall frequency response follows that of the radar response with no wall intervening. The response of the radar is convolved with the actual frequency response of the wall.

In order to remove the frequency response of the radar it must be deconvolved from the wall plus radar response data. In the frequency domain this is performed by dividing the frequency response of the combination by the frequency response of the radar alone. The result is the normalized frequency response of the wall. By ensuring the gain settings of the radar for both the antenna to antenna radar response measurement and the through transmission wall measurement are the same, the normalization process results in a curve of the absolute response of the wall versus frequency.

For each of the walls, we took the antenna to antenna, no wall radar measurements and the through wall transmission measurements and placed a 10 nsec window around the peak response. This allowed us to reduce the effects of multipath on the result. The deconvolution process was applied and the resulting wall response versus frequency was plotted. Figure 40 shows a comparison of the frequency response of each of the walls plotted on the same graph. As can be seen from the graph drywall is virtually transparent, brick shows some attenuation at the higher frequencies, the response of adobe and block walls is very similar, and concrete shows the most attenuation of all.

The bulk dielectric response of each wall was also determined by using the static test data. In this case we used the transit time difference between the through wall transmission measurement and the antenna to

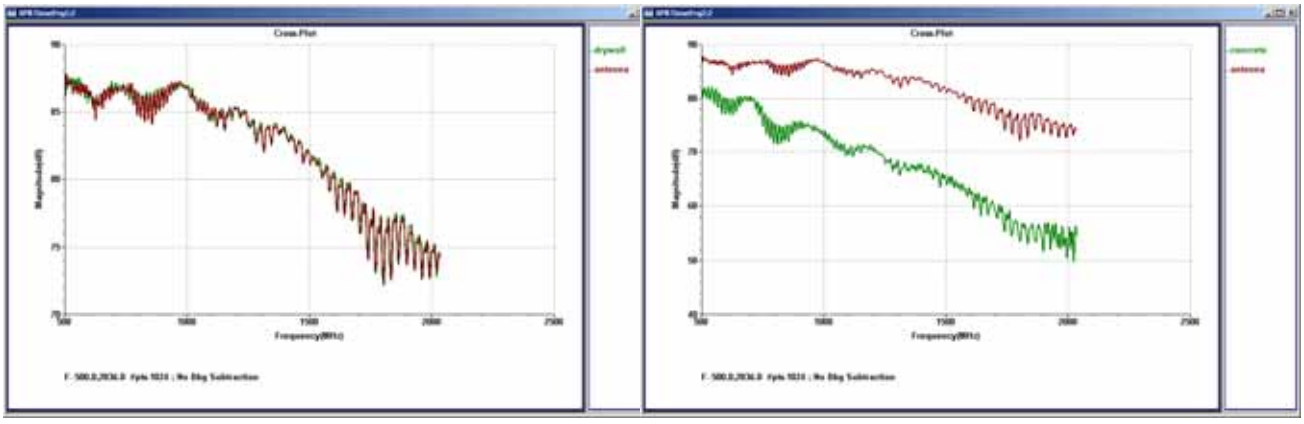

Figure 39 - The frequency response of a wall contains the frequency response of the radar.

antenna radar measurement plus the thickness of the wall to calculate the dielectric constant. As shown in Equation (1)

$$\varepsilon = (c*\Delta t/ d_{wall} + 1)^2 \qquad (1)$$

the dielectric constant is calculated from the time difference of arrival between the wall and no wall case, the thickness of the wall, and the speed of light. Figure 41 shows a graph comparing the transit time difference between air and concrete for the wall.

Equation 1 was used to calculate the average dielectric constant of each material over the entire frequency band. Table 2 shows the results of the calculations. Additional calculations using portions of the frequency band showed that there is only a weak dependence with frequency so the average value for the band of frequencies tested is a good approximation.

Motion detection performance was determined by reading the motion detection data collected into our Open GL radar control display program and using its features to change the motion detection algorithms as the image reconstruction proceeded. As a result of our testing we had 18 different sets of data to use for verification purposes.

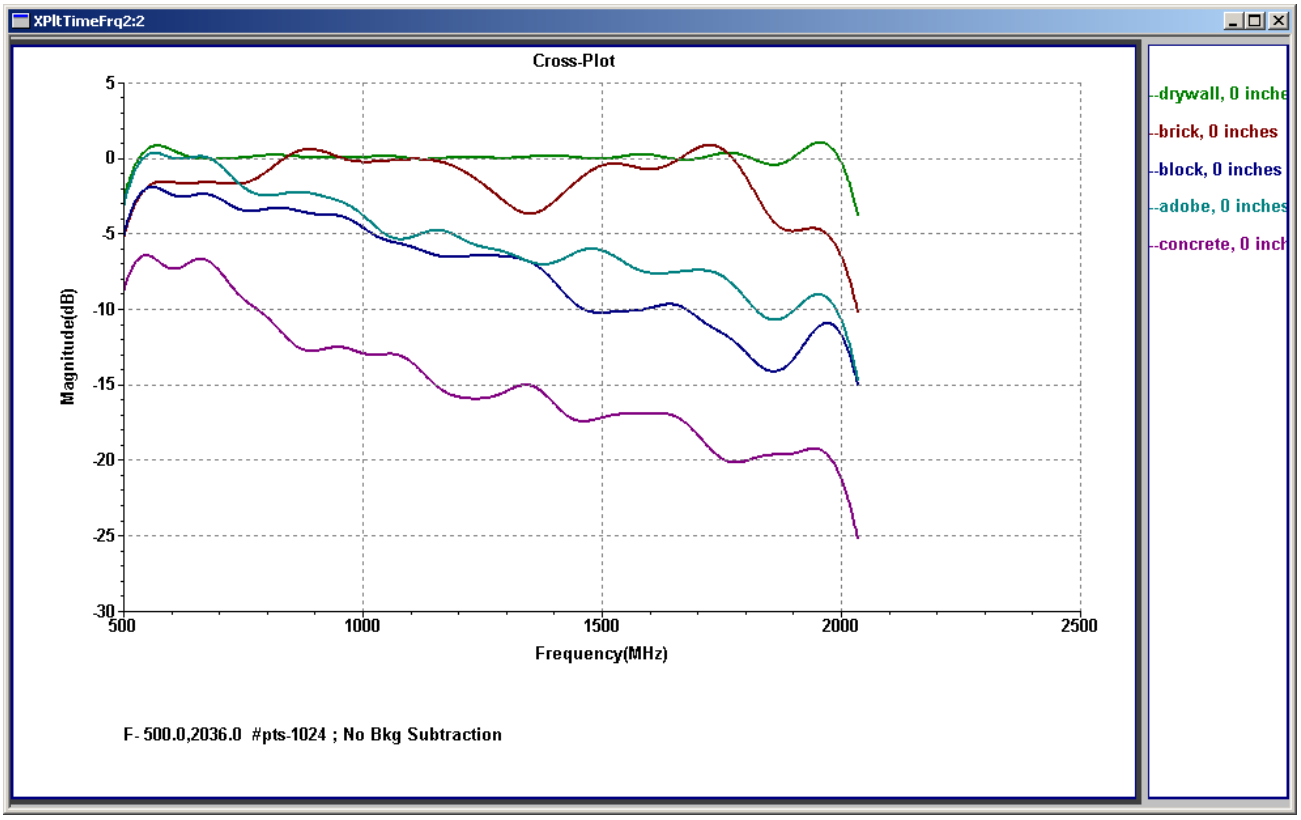

Figure 40 - Attenuation characteristics of wall materials.

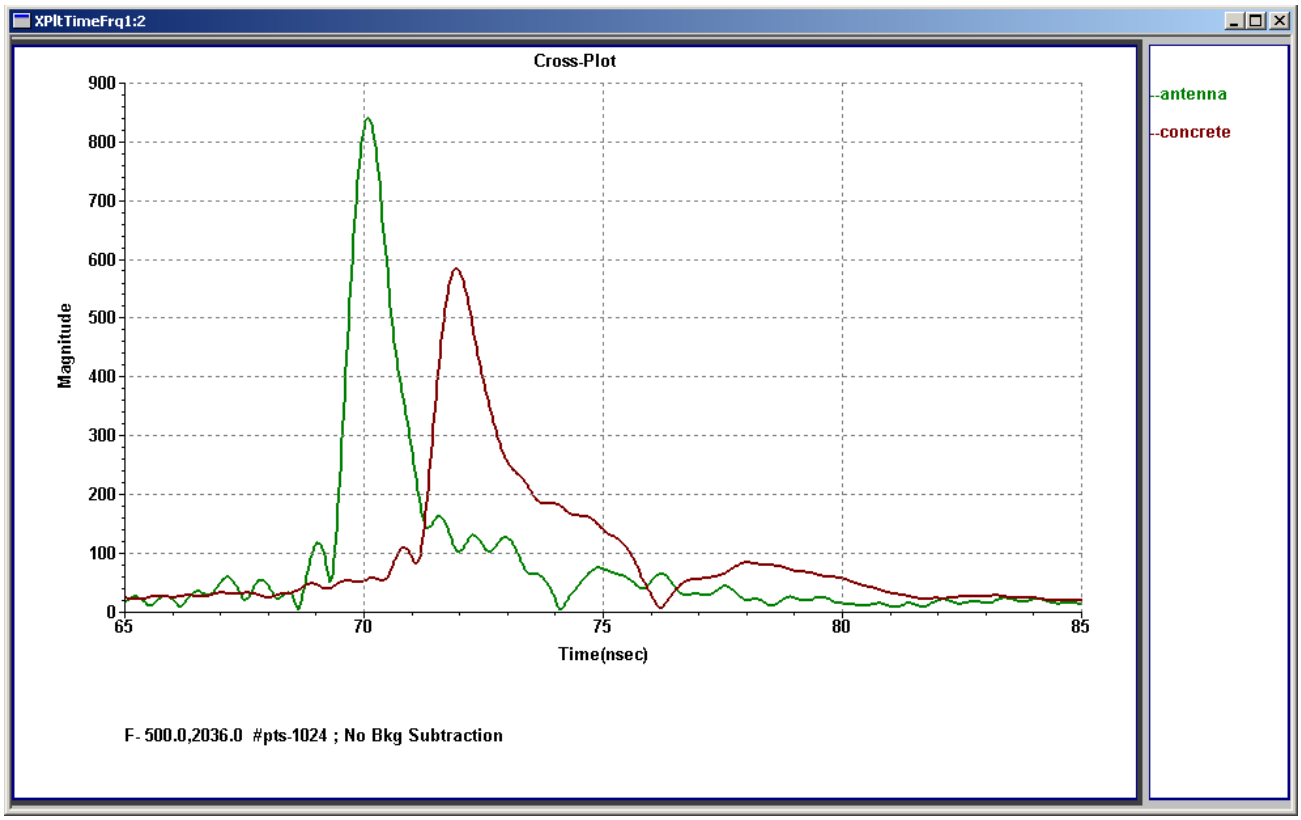

Figure 41 - The transit time through concrete is greater than through air.

Table 3 summarizes the results of the verification testing for the data taken at Ft. Monmouth. As shown in the table we were able to detect both large and small motions through each of the test walls from significant standoff distances. Figure 42 shows a frame from concrete motion detection verification test where a person was sitting in a chair 12' behind the concrete wall with the antennas at a standoff distance of 10 feet. Even though the person was sitting as still as possible they are still clearly detected. These results were

Drywall - antennas at wall	$\varepsilon \approx 1.04$
Brick - antennas at wall	$\varepsilon \approx 2.43$
Brick - antennas 30" from wall	$\varepsilon \approx 1.95$
Adobe - antennas at wall	$\varepsilon \approx 4.52$
Block - antennas at wall	$\varepsilon \approx 4.48$
Block - antennas 27" from wall	$\varepsilon \approx 4.12$
Concrete - antennas at wall	$\varepsilon \approx 7.85$
Concrete - antennas 28" from wall	$\varepsilon \approx 6.83$

Table 2 - Average dielectric constant of wall materials.

Drywall: at 19´ from standoff of 12´
Brick: at 21´ from standoff of 14´
Concrete block: at 28´ from standoff of 21´
Adobe: at 16´ from standoff of 4´
Reinforced concrete: at 22´ from standoff of 10´

Table 3 - Distances of detection for breathing and large motion for test walls.

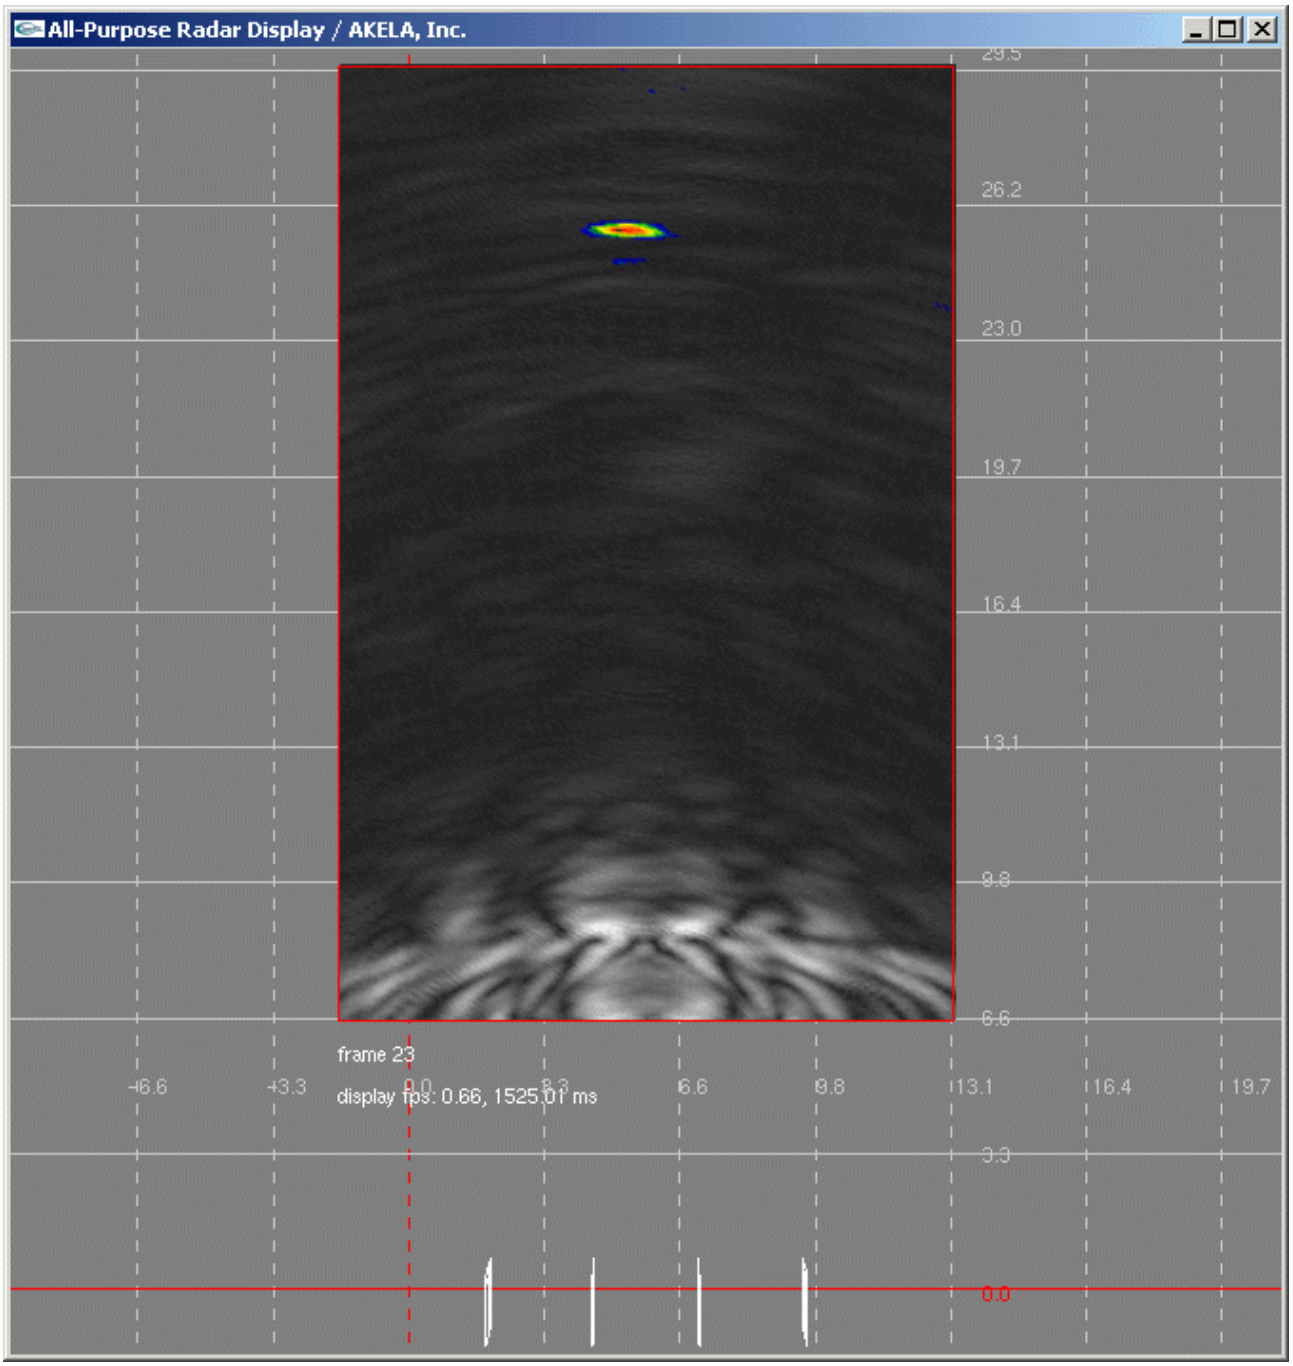

Figure 42 - A person sitting 12' behind a concrete wall.

significant since they showed that the sensitivity of the radar and motion detection algorithms was sufficient to detect breathing at significant standoff distances through steel reinforced concrete.

Additional testing was performed in Santa Barbara to gather static and dynamic data for determining radar detectability, motion algorithm sensitivity, and to get material property data on stucco and reinforced concrete walls. As with the testing at Ft. Monmouth, we used the two antenna radar configuration for the static testing and the eight antenna configuration for the dynamic testing. To make data collection easier, we used two different fixed antenna configurations. One collapsible configuration consisted of four antennas spaced equally over a span of 2.9 meters (8'4") and the other over a span of 4.9 meters (16'). Figure 43 shows these different configurations.

Figures 44 and 45 show the results of the wall testing. In Figure 44 are two graphs of the attenuation versus frequency for steel reinforced concrete. On the left in the figure is the result obtained from the wall tested at Ft. Monmouth. On the right is the result of a similar measurement made on a 12" reinforced concrete wall from a building at the University of California, Santa Barbara campus. As shown in the figure, the attenuation is virtually identical to the Ft. Monmouth test wall. This suggests walls made with standard construction techniques should exhibit similar attenuation response.

Static tests on two different stucco walls were also performed in Santa Barbara. The stucco in each wall consisted of chicken wire over a tar paper base on top of which was a 1" stucco finish. One wall had a standard drywall face on the interior portion of the wall while the other was a double sided stucco wall. Figure 45 show the comparison of the attenuation response of these two walls. As can be seen in the figure, stucco walls are relatively transparent above about 900 MHz and exhibit slightly less attenuation than a block wall. However, especially in the case of the double sided wall, the low frequencies exhibit too much attenuation to be useful.

Static testing in support of our simulation activities consisted of making measurements of a person at various distances from the radar in an open field test environment. The goal of these tests was to reduce the amount of environmental noise present in the data so that we could get a best case estimate of the range limits of the radar under low clutter conditions. Tests using the two antenna configuration were conducted at distances between 5 and 30 meters at 5 meter intervals. Tests using a four element array were performed to determine the additional range benefit gained by static imaging. These tests were performed with different radar sample rates and different antenna separations at distances out to 100 meters.

Because we were unable to find a suitable test site in Santa Barbara where we could achieve significant standoff distances to walls of various types, our motion testing was confined to an open field environment. Tests were performed with both antenna arrays over a range of radar sample rates in order to get data for determining the maximum range at which we could see motion using our motion detection algorithms. Using

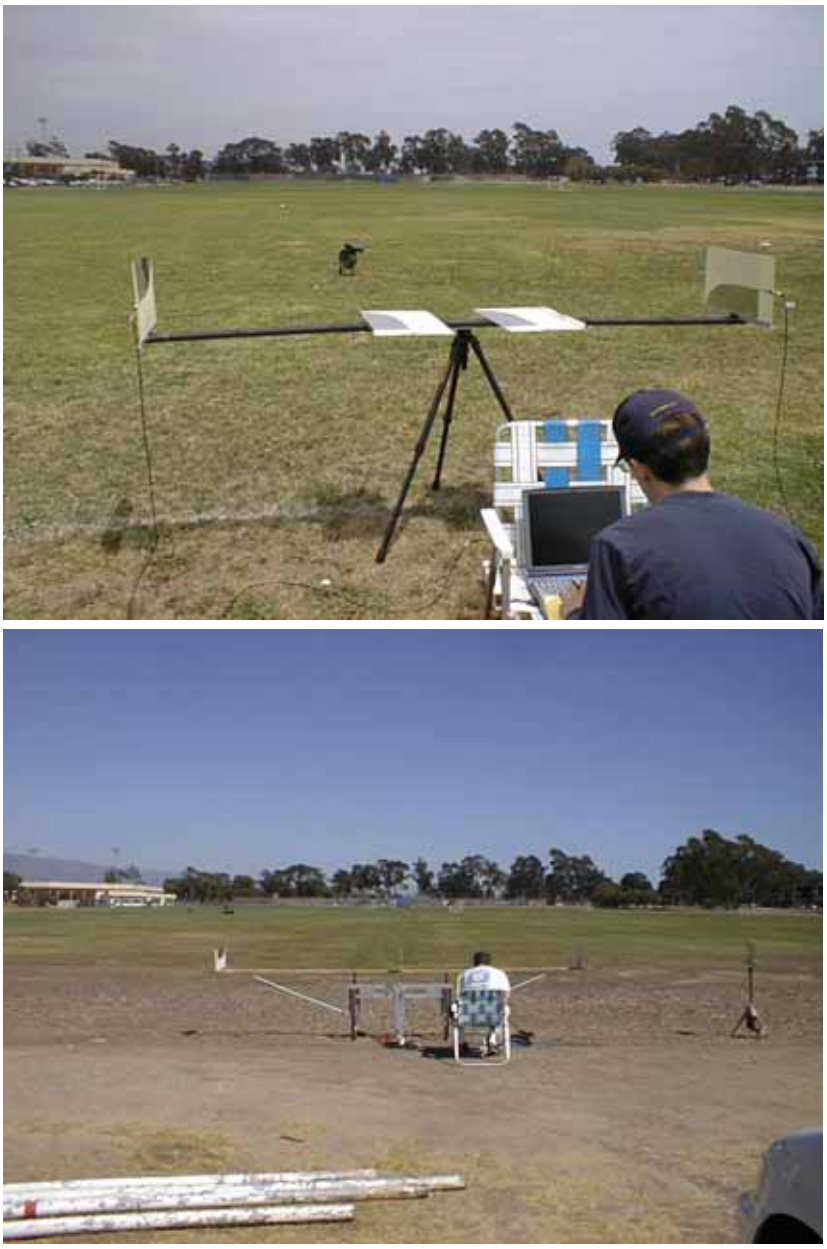

Figure 43 - Antenna array setups for static and dynamic testing.

this maximum range data with the attenuation characteristics of the walls we measured enables us to determine the maximum likely motion detection range of our imaging radar through the various wall materials.

The open field static test data was used to determine the range limit of our radar hardware. To analyze this data, we Fourier transformed each data set to get a time domain range profile. The target was located in each data trace and its magnitude compared to the magnitude of the quiet background. This allowed us to calculate a signal to noise ratio for the target return at each distance. The log of the signal to noise ratio was plotted against distance and compared to the level of the environmental background. Figure 46 shows plots of the data for a person for both front and side aspect angles. For distances beyond 30 feet, the slope of

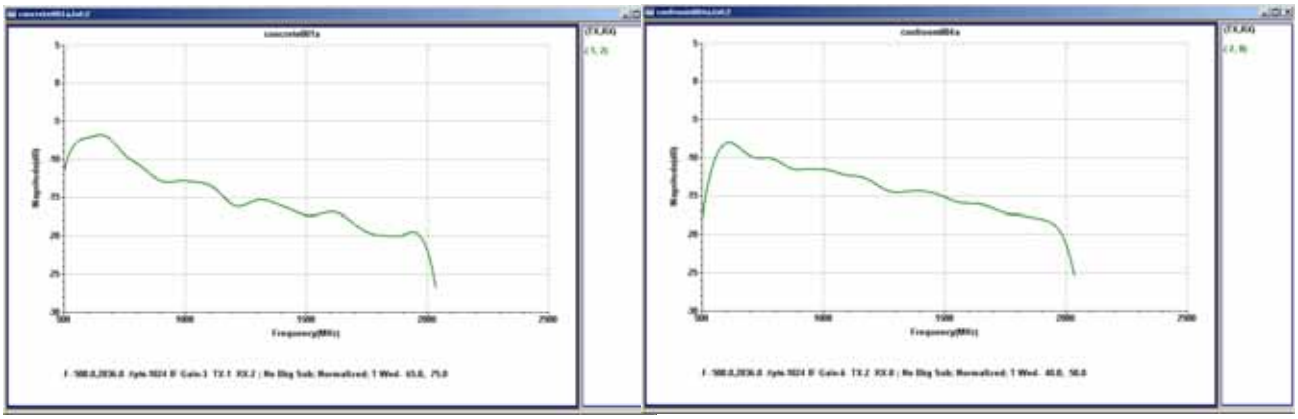

Figure 44 - Comparison of attenuation response of two 12" reinforced concrete walls. left: Ft. Monmouth, right: UCSB

the line fitting the data points is inversely proportional to range to the fourth power, which the expected result for far field scattering. Extrapolating the line until it intersects the noise floor gives the theoretical distance at which the target return is indistinguishable from the noise. As shown in the figure, this value is between 250 and 300 feet.

This data was also used to determine the sensitivity of the motion detection algorithms. In this case, the static data was modified in a MATLAB program to simulate movement. Each data set was offset by a set

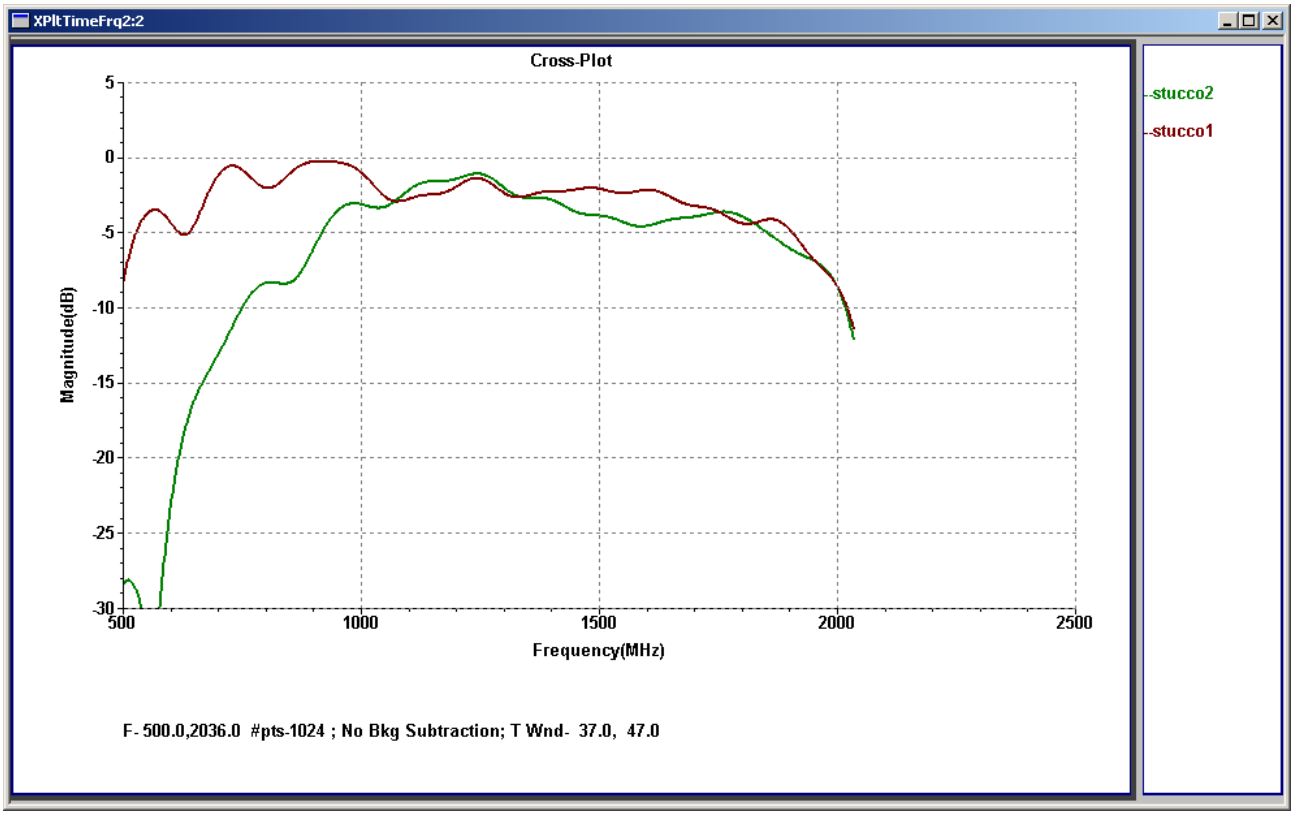

Figure 45 - The attenuation response of single and double sided stucco walls.

of predetermined distances to create a series of data frames that were then used in the motion detection algorithms to show the way in which signal to noise ratio changes with the degree of movement for each algorithm. Figure 47 shows the results of these simulations. As shown in the figure, the small motion algorithm is the most sensitive of all. The first order motion detection algorithm is more sensitive to small motions than the second order algorithm.

Since the sensitivity of the algorithm to movement is a function of signal to noise ratio, the amount of motion that is detectable for any imaging situation will be determined by target size, distance, and wall thickness. This makes sensitivity of the imaging radar a very important parameter. For standoff ranges that are short and walls that are relatively transparent, even the large motion algorithms have the ability to detect small motions. We have seen this behavior in some of the motion detection data analysis we have done on the data through the test walls.

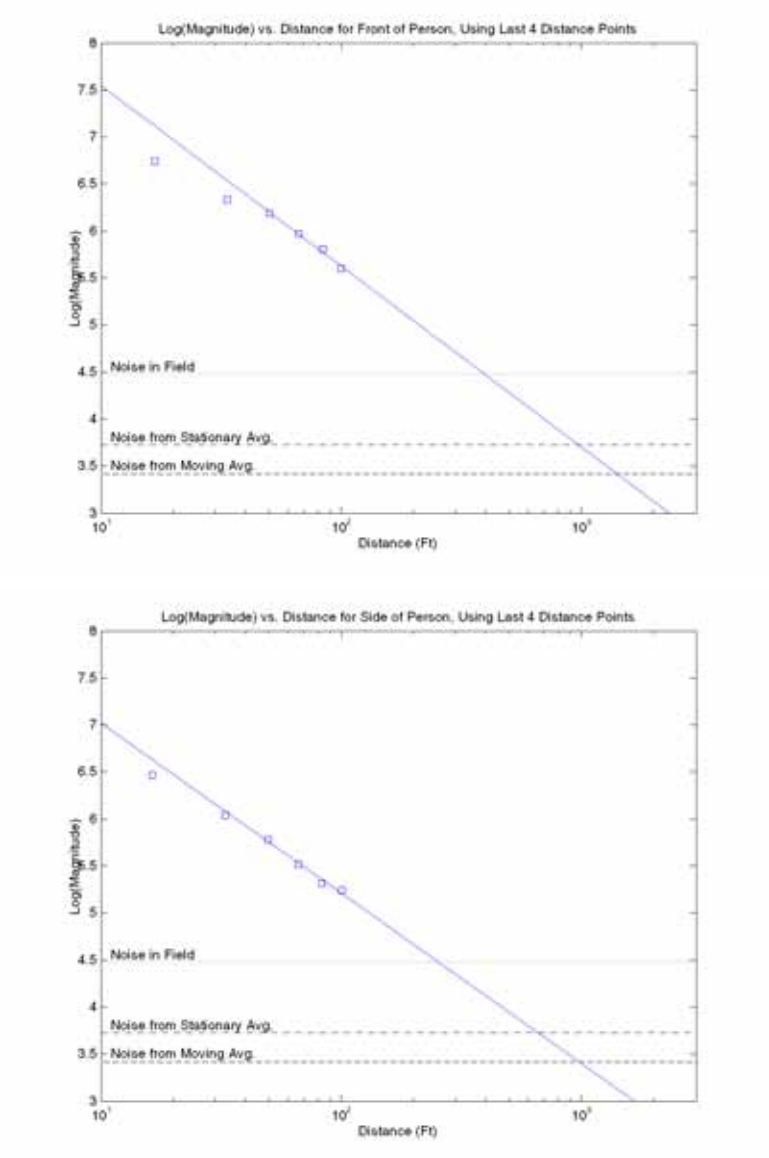

Figure 46 - Detection limit of radar determined from experimental data.

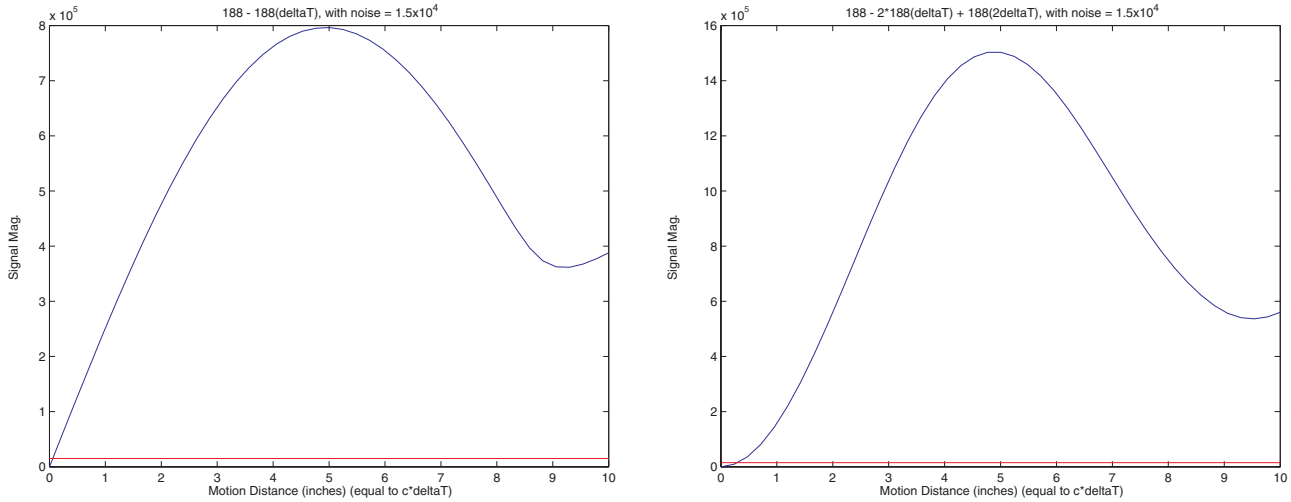

First order

Second order

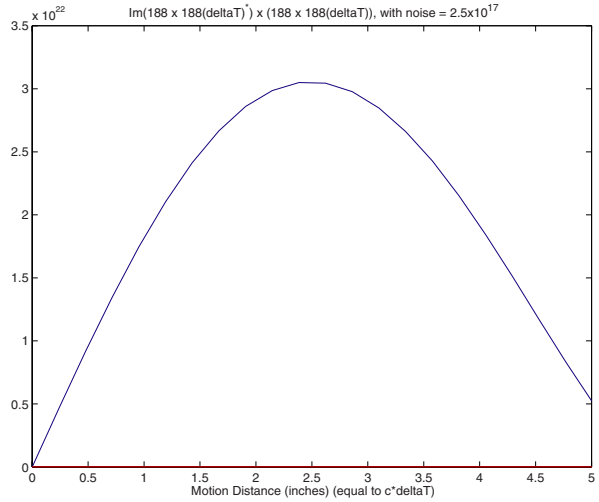

Small motion

Figure 47 - Motion detection algorithm sensitivity.

Analysis of the motion data taken in the open field at Santa Barbara was performed by reading the motion detection data collected into our Open GL radar control display program and using its features to change the motion detection algorithms as the image reconstruction proceeded. On the basis of our earlier static test analysis our expectation was that we would be able to see motion out to a distance of 300 feet. Achieving this range would be equivalent to seeing motion through a 12" thick concrete wall at a distance of 100 feet. Unfortunately, the results showed that we could see motion reliably out to only about 200 feet as shown in Figure 48.

Subsequent analysis pointed to several reasons for the shortfall. The first was the difference in sample rate. While taking data at the highest sample rate of the radar is advantageous from the point of view of capturing fast movement, the noise the radar adds to the data also varies linearly with the sample rate. Since the sample rate of the open field tests was three times that of the tests at Ft. Monmouth, the subsequent data is

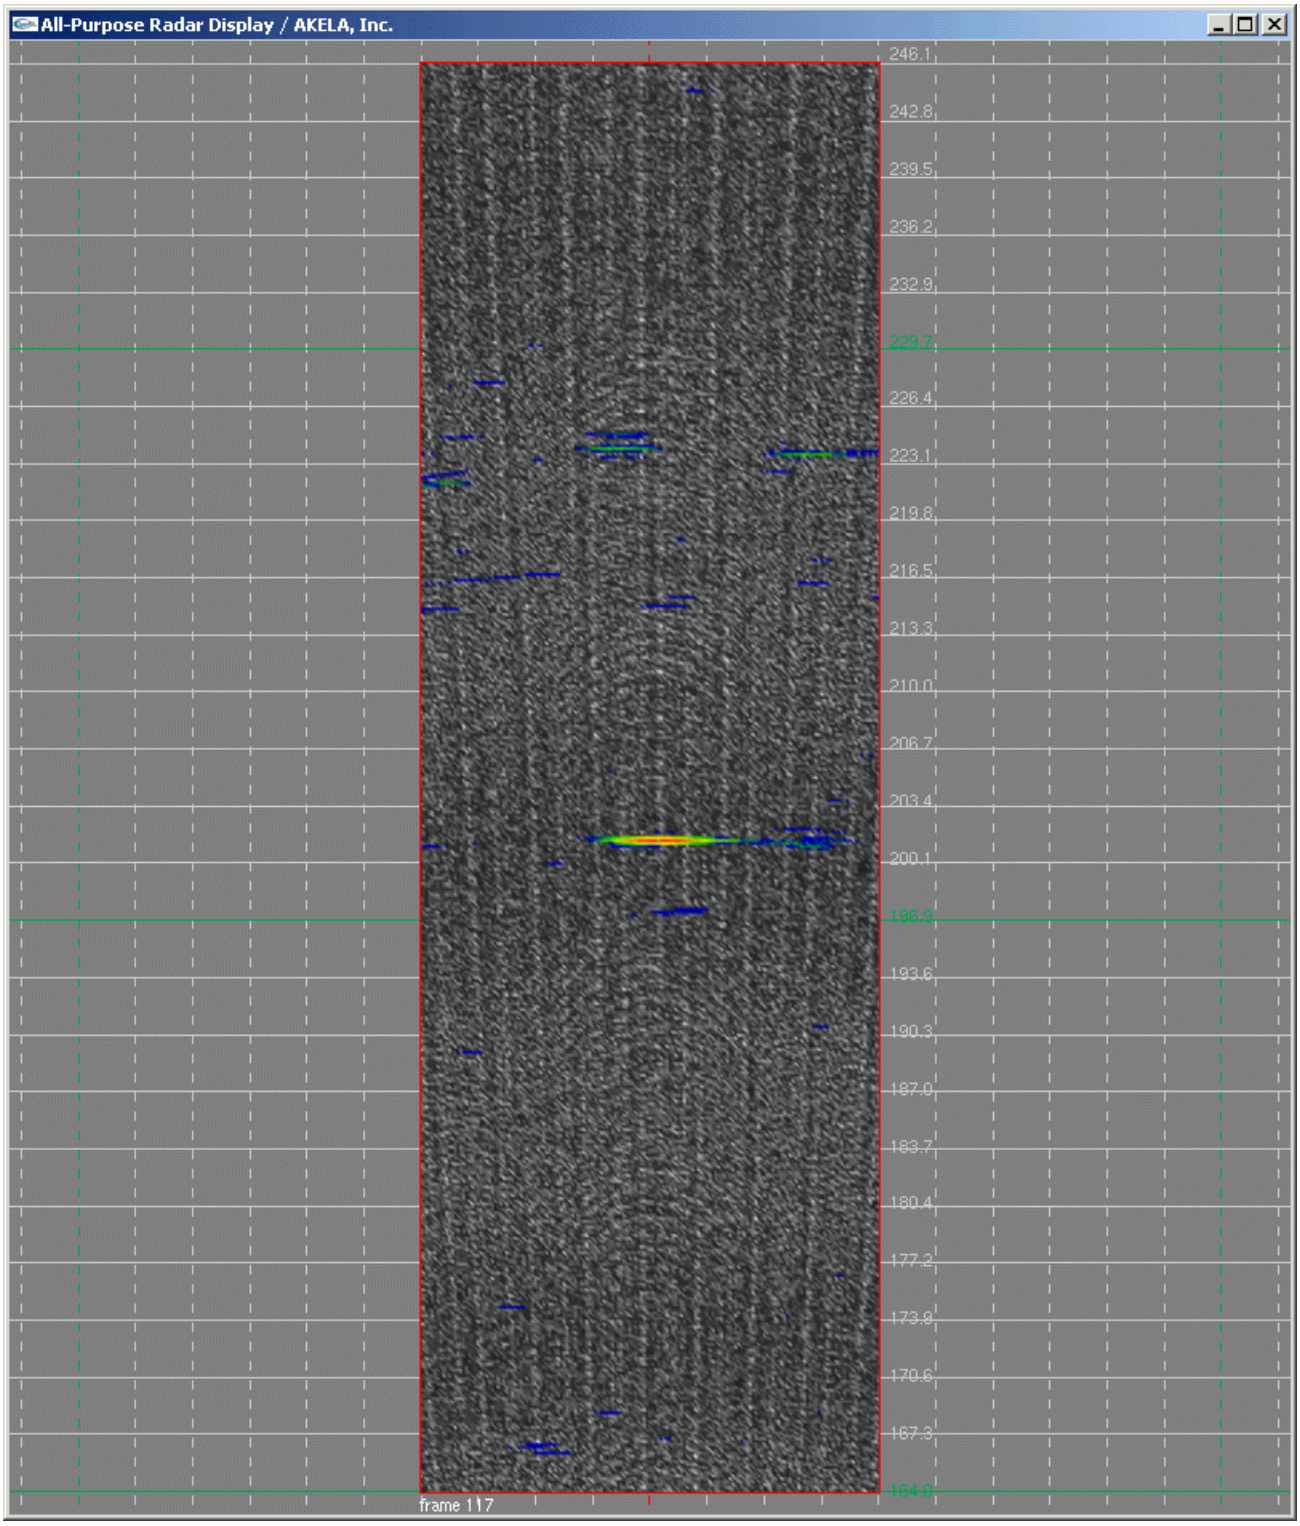

Figure 48 - Image of motion detected at 200 feet in open field.

much noisier so the ultimate range performance is about the same for the two sets of data. As indicated earlier, a 200 foot range in the open field is roughly equivalent to a 50 foot range through a 12" thick concrete wall. Internal sources of noise in the radar contributed to the poorer than anticipated performance as well.

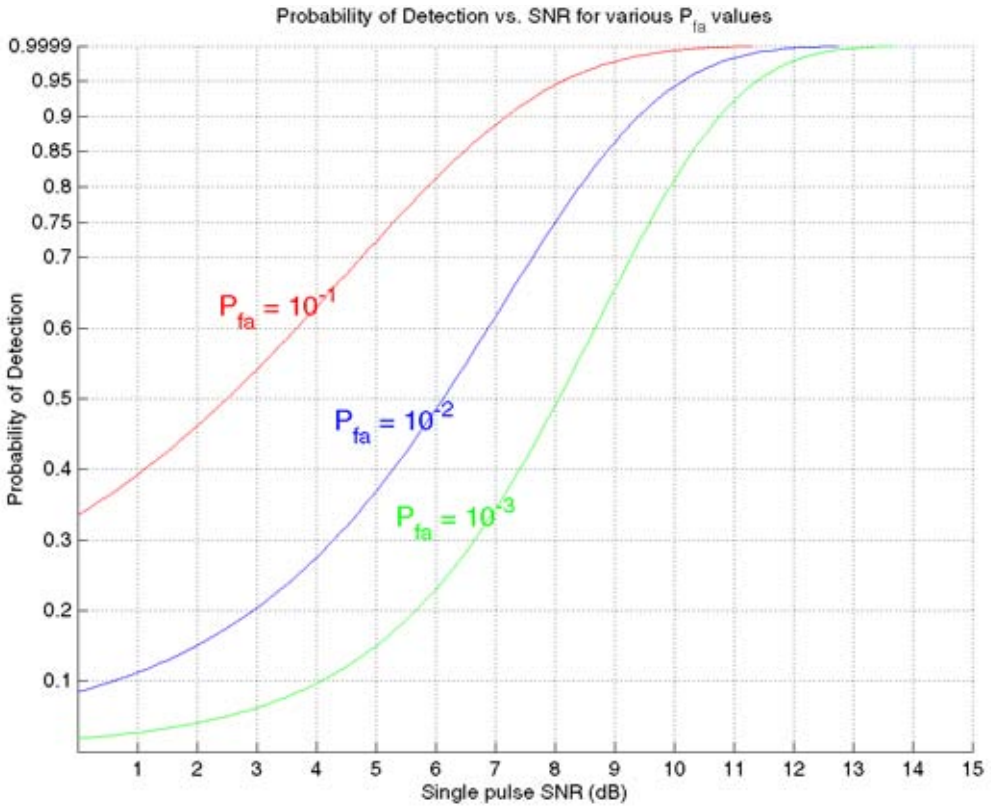

Figure 49 - Single sweep probability of detection.

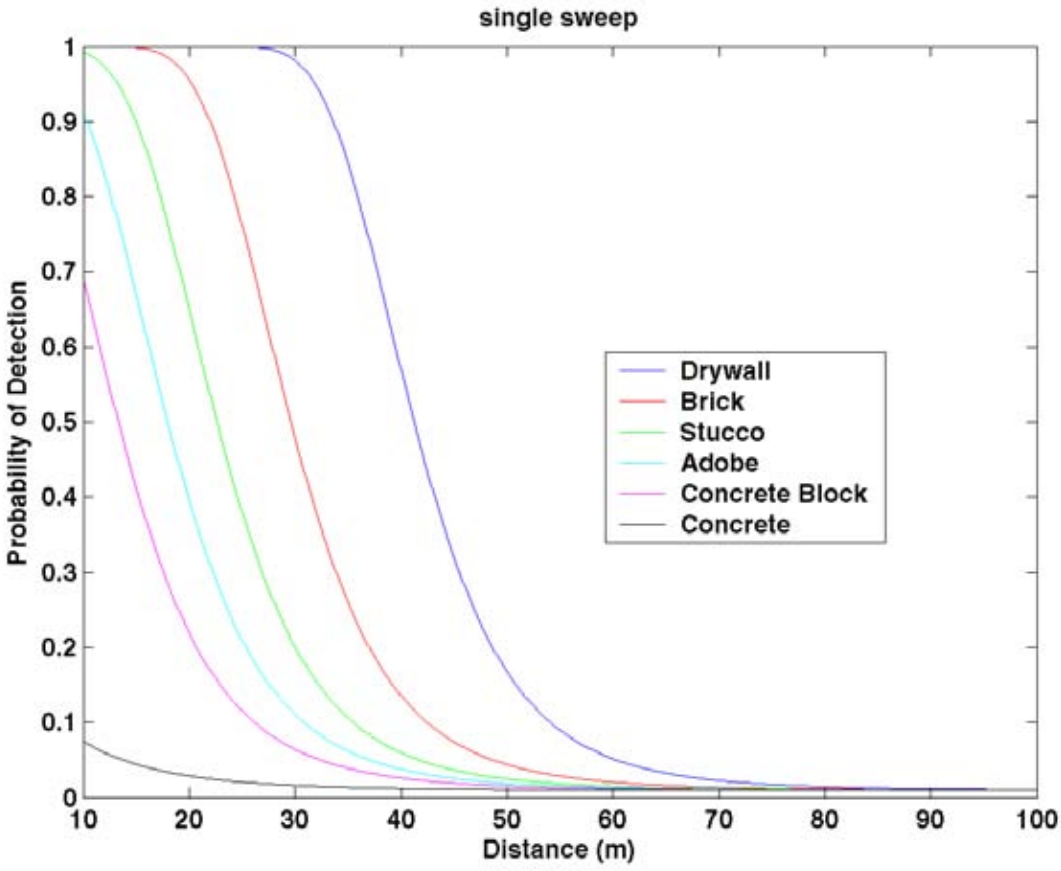

Figure 50 - Single sweep probability of detection for person through various wall materials.

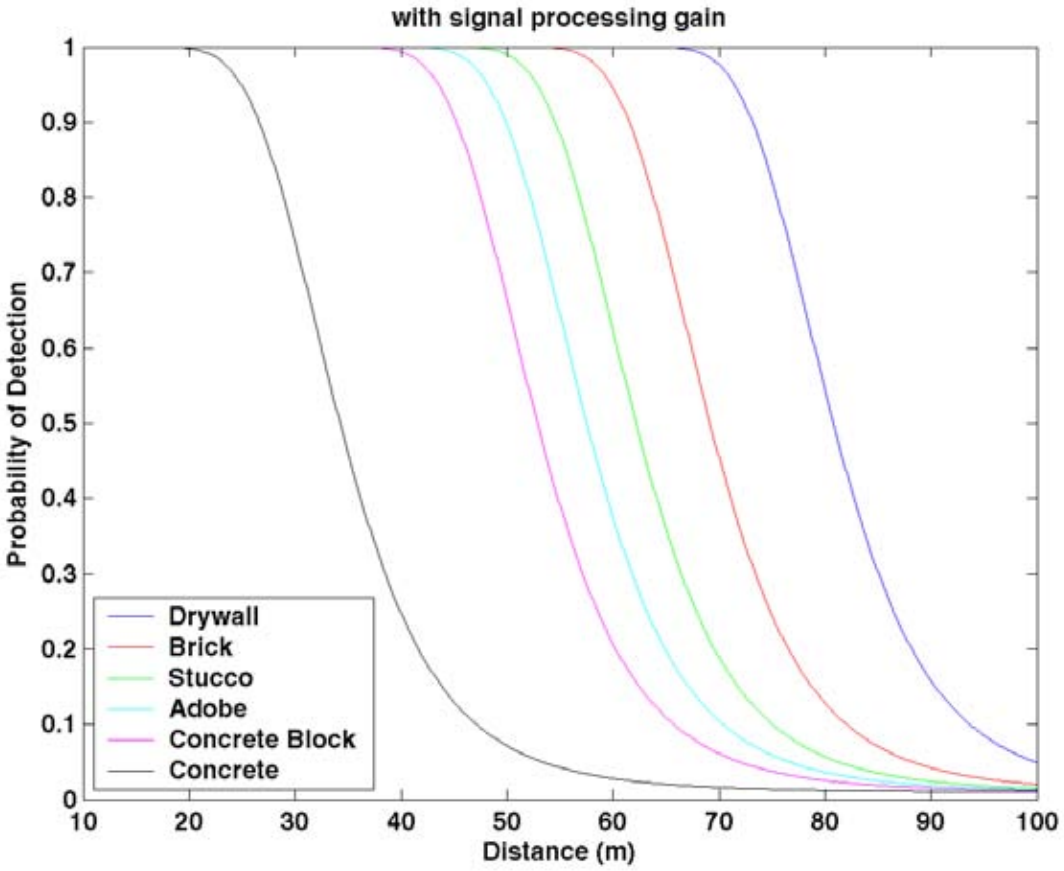

Figure 51 - Probability of detection through various wall materials using motion detection.

We decided to perform a more rigorous analysis using the static data to make predictions of the probability of detecting motion through the various wall types. Detection of signals in noise is a well studied problem in the radar literature. One method of calculating the probability of detecting a signal as a function of noise is Parl's method. Figure 49 shows a plot of the probability of detection versus signal to noise ratio for a set of fixed false alarm rates calculated for a single radar sweep using this method.

Our experimental data allows us to plot signal to noise ratio as a function of distance for a person. We used a linear fit of the data to generate a smooth curve of probability of detection versus distance. By adding in the attenuation due to a wall, it is possible to generate the single sweep probability of detecting a person through that wall. Figure 50 shows the result for the walls we tested at Ft. Monmouth. The calculations performed are for a false alarm rate of 1%.

Signal processing adds gain to the system since it uses multiple sweeps. Figure 51 shows the probability of detection versus distance for the same conditions but adding in the extra gain provided by first order motion detection. As can be seen in the figure, there is considerable improvement. These predictions are consistent with the open field data of Figure 48.

Testing was also performed in support of location algorithm development. Image reconstruction using high resolution range profiles relies on knowing the location of both the transmitting and receiving sensors. Movement of either sensor translates the range profile along the time (range) axis. The result is subimage (the image created from a single sensor pair) misregistration and subsequent blurring of the overall image. The operating characteristics of our sensor system provide information that is used to determine sensor location directly from the data.

The direct path signal between the transmitter and receiver of the sensor pair will always be the first and largest peak in the range profile. Determining where the peak occurs in time and then using the speed of light gives the absolute distance between the transmit and receive antennas. This distance is measured for each of the antenna pairs and then the positions of each of the sensors is solved for using techniques widely used for emitter location.

The requirement for any of these techniques is that there be one more observation than the number of dimensions for which the position is desired. In our case, we wish to know the position in 2 dimensions since our program objectives did not include situations where there is vertical antenna aperture. For a 4 sensor system all of the combinations of transmitters and receivers give us $(N_t) \times (N_r-1) = 12$ observations which is greater than the minimum needed to locate all 4 sensors in two dimensions.

The accuracy of position location depends on the distance measurements for each sensor pair. Our sensor system bandwidth sets the fundamental distance resolution at 0.1 meter. However, it is possible to use the Fourier transformation process to accurately interpolate between data points. This is accomplished by selecting a larger length FFT, placing the sensor data at the correct starting position in the longer input data set, and then filling with zeros on either end before doing the transformation. For example, a 512 point data set using an FFT length of 8192 points gives an effective distance resolution between range profile points of $0.1/16 = 0.00625$ m or about 0.25 in. Figure 52 shows an example of this technique for providing "super resolution" comparing a 1024 and 8192 point FFT of the same data. Note that there is a shift in position of the peak of the waveform. Using a 16384 point FFT would theoretically improve the resolution by an additional factor of two.

For radars that have independent clocks, finding the peak is complicated by the phase variations in the different clocks as indicated earlier. We used independent units to test the peak finding and location algorithms. Figure 53 shows photographs of a single independent radar sensor consisting of the radar, interface board, discone antenna, and wireless radio. As can be seen in the figure, there is a provision for using a wired Ethernet interface that was used extensively to improve robustness of communications during development testing.

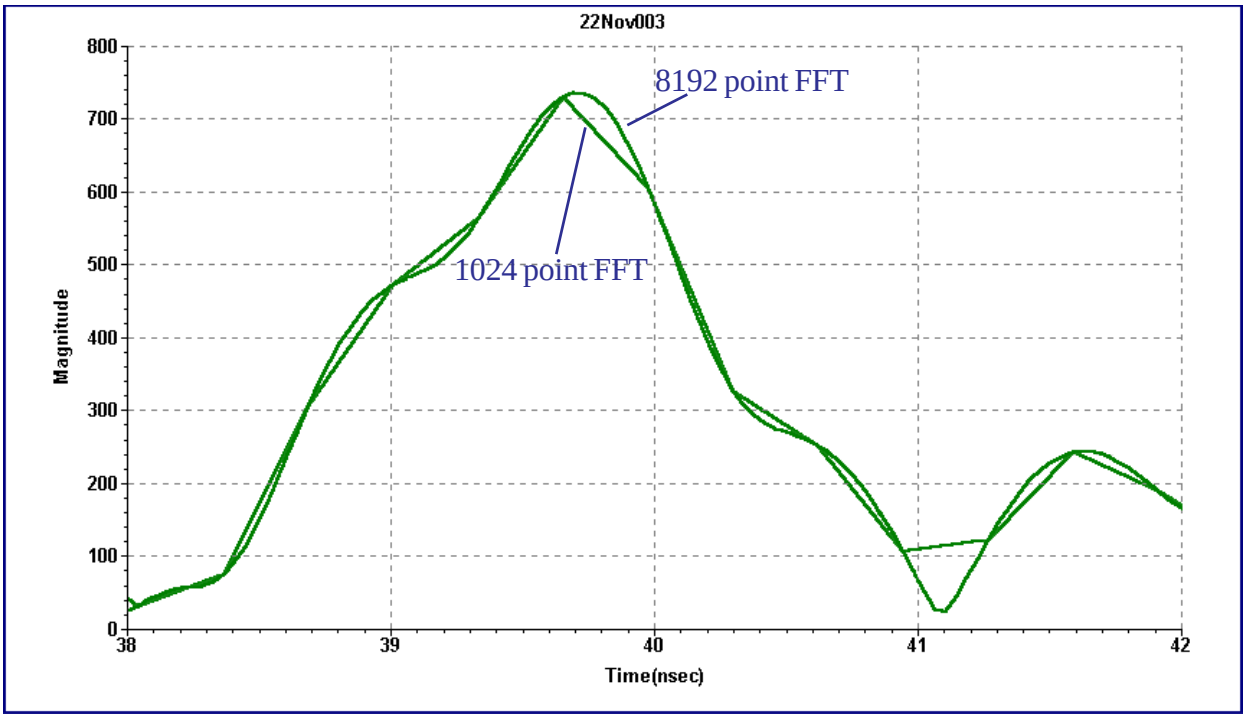

Figure 52 - FFT interpolation can improve distance measurement.

Figure 53 - Independent radar sensor units.

Since the clocks are not initially aligned, the radar control software is first used to perform a clock synchronization. Once the radar clocks are synchronized, the radars will communicate and sweep data can be Fourier transformed to obtain the range profiles between radar sensor units. However, there is a residual uncertainty in the independent clocks that must be removed before the location algorithm can be used. If the clock in the transmitting sensor is slightly ahead of the clock in the receiving sensor, it will appear to the receiving sensor that the transmitter is closer than it actually is. If the receiving clock is faster than the transmitter clock, it will appear that the transmitter is farther away than it actually is. To find the correct distance between the transmitter and receiver, it is necessary to do bidirectional scans. Comparing the distance differences between the scans when the sensors switch roles allows us to correct for the clock mismatch and find the actual distance between the sensors.

Figures 54 and 55 show data traces from an array of four independent radar units. Figure 54 shows the result of performing a complete set of sweeps on a set of four independent radar units after the initial clock synchronization has been completed. The graph at the bottom of the figure is the frequency sweep data from each unit and the top trace is the Fourier transform of the frequency data. The two sided Fourier transform is plotted here since it is possible after initial clock synchronization for a trace to show up in negative time. As can be seen in the figure, the peaks are distributed widely. For example, the separation between unit 6 and unit 8 is of the order of 50 nsec or about 50 feet where in actuality the units were less than 4 feet apart. This difference is due to the clock misalignment.

The data from each pair of scans is used to perform the alignment correction. In order to find the correct distance between each pair of sensors, we use a cross correlation technique. This consists of multiplying the frequency data from the first scan with the complex conjugate of the frequency data of the complementary bidirectional scan, and then taking the Fourier transform. The peak in the resulting transform is at twice the time difference between the two units. This value is then used to phase adjust the original frequency data before further processing is done.

Figure 55 show the results of the alignment process. The phase adjusted frequency data are shown in the bottom graph of the figure and the Fourier transform of the data is shown in the top graph. As can be seen in the graph, the peaks of all of the range profiles are clustered and spread over several nanoseconds which represents the actual separation between the independent units.

Once the alignment process has been completed, the distance between the units can be determined by the same correlation process as for the original alignment. We have been able to demonstrate that the corrected independent unit data can be used to determine the locations of the individual sensors and have incorporated the algorithm into the system control and image processing software.

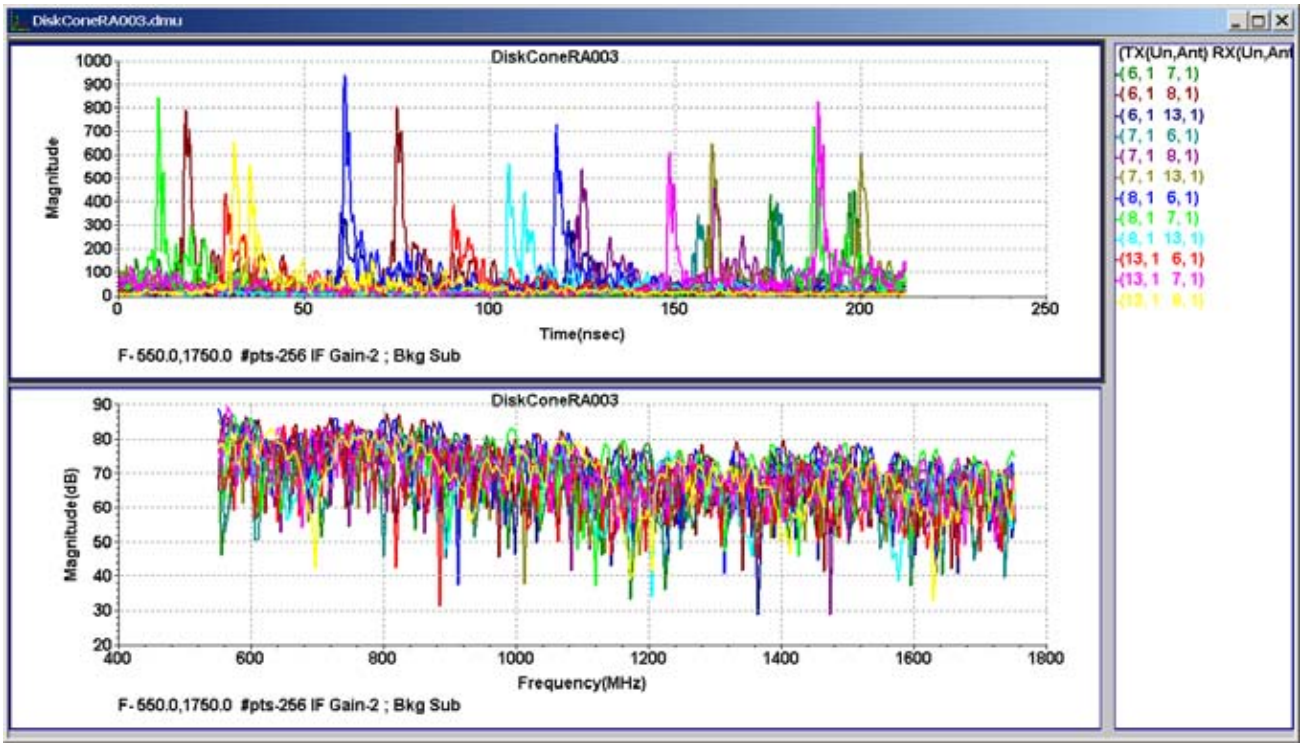

Figure 54 - Independent unit data traces before clock correction.

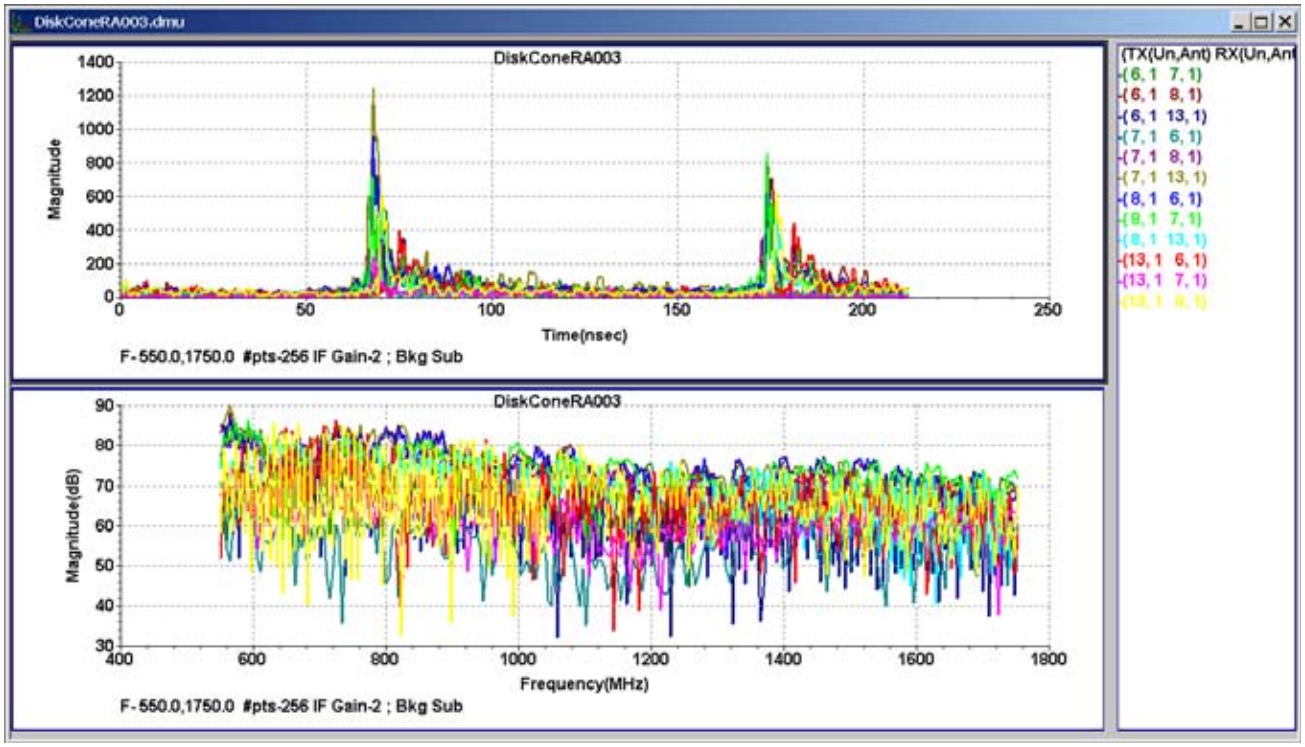

Figure 55 - Independent unit data traces after clock correction.

The image processing software works with the independent units but the results have not been as good as with the single radar fixed array system. Objects appear in the image but are not as clearly defined. The reasons for this appear to be that there is more noise in the independent unit imaging system since the noise from the individual clocks are uncorrelated, and that there is a larger impact on frequency drift than we anticipated. To improve imaging performance will require additional effort on the hardware side to reduce the radar noise, and on the software side to implement more complex data correction and image processing methods.

Conclusion

Virtually all of our original program goals have been met. We have developed a small, robust, cost effective radar. We have demonstrated the capability to detect both large and small motions with the radar at substantial standoff distances through steel reinforced concrete walls. The software we have developed to control the radar and perform imaging has demonstrated the capability to perform imaging in real time at frame rates up to 5 frames a second. We have characterized in detail the frequency dependent attenuation of most common wall types. The system has demonstrated the ability to be used with antennas in random locations using wireless control. We have made detailed measurements and predictions of the system's probability of detecting individuals through different types of walls as a function of distance and shown that it is possible to detect a person through 12" of steel reinforced concrete at a range of 20 meters. And we have demonstrated the ability for sensor self location. While imaging with the random sensor system was not as good as anticipated, we understand what must be changed to improve system performance.